

forall x

Adelaide 2017

P.D. Magnus

University at Albany, State University of New York

Tim Button

University of Cambridge

Antony Eagle

University of Adelaide

© 2005–2017 by P.D. Magnus, Tim Button, and Antony Eagle. Some rights reserved.

This book is a derivative work created by Antony Eagle, based upon Tim Button’s Cambridge version of P.D. Magnus’s `forallx` (version 1.29). There are pervasive changes, big and small. Button’s version is available at www.nottub.com/forallx.shtml; Magnus’ original is available at fecundity.com/logic. Both texts were released under a Creative Commons license (Attribution-ShareAlike 3.0), as is the present text.



You are free to copy this book, to distribute it, to display it, and to make derivative works, under the following conditions: (a) Attribution. You must give the original author credit. (b) 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 creativecommons.org/licenses/by-sa/3.0/

This version of `forallx`-Adelaide is current as of July 25, 2017.

Typesetting was carried out entirely in $\text{X}_{\text{Y}}\text{L}_{\text{A}}\text{T}_{\text{E}}\text{X}$. The body text is set in Constantia; math is set in Cambria Math; sans serif in Calibri; and monospaced text in Consolas. The style for typesetting proofs is based on `fitch.sty` (v0.4) by Peter Selinger, University of Ottawa.

Contents

1	Key notions	1
<hr/>		
1	Arguments	2
2	Valid arguments	4
3	Other logical notions	10
2	Sentential logic	14
<hr/>		
4	First steps to symbolisation	15
5	Connectives	19
6	Sentences of Sentential	30
7	Use and mention	35
3	Truth tables	40
<hr/>		
8	Schematic truth tables	41
9	Truth-functional connectives	44
10	Complete truth tables	48
11	Semantic concepts	54
12	Truth table shortcuts	61
13	Partial truth tables	66
14	Expressiveness of Sentential	70
4	Quantified logic	74
<hr/>		
15	Building blocks of Quantifier	75
16	Sentences with one quantifier	82
17	Multiple generality	93
18	Identity	101
19	Definite descriptions	107
20	Sentences of Quantifier	115

5 Interpretations	120
<hr/>	
21 Extensionality	121
22 Truth in Quantifier	127
23 Semantic concepts	137
24 Demonstrating Consistency and Invalidity	140
25 Reasoning about all interpretations: demonstrating inconsistency and validity	146
6 Natural deduction for Sentential	150
<hr/>	
26 The very idea of natural deduction	151
27 Basic rules for Sentential	154
28 Proof-theoretic concepts	180
29 Proof strategies	183
30 Derived rules for Sentential	185
7 Natural deduction for Quantifier	193
<hr/>	
31 Basic rules for Quantifier	194
32 Derived rules in Quantifier: conversion of quantifiers	208
33 Rules for identity	211
34 Proof-theoretic concepts and semantic concepts	216
Appendices	221
<hr/>	
A Alternative terminology and notation	221
B Quick reference	224
Acknowledgements, etc.	229

Chapter 1

Key notions

1 Arguments

Logic is the business of evaluating arguments – identifying some of the good ones and explaining why they are good.

In everyday language, we sometimes use the word ‘argument’ to talk about belligerent shouting matches. Logic is not concerned with such teeth-gnashing and hair-pulling. They are not arguments, in our sense; they are disagreements.

Giving an argument, in the sense relevant to logic (and other disciplines, like law and philosophy), is something more like *making a case*. An argument presents reasons that purport to favour – or support – a specific claim. Consider this example:

It is raining heavily.
If you do not take an umbrella, you will get soaked.
So: You should take an umbrella.

We here have a series of sentences. The word ‘So’ on the third line indicates that the final sentence expresses the CONCLUSION of the argument. The two sentences before that express PREMISES of the argument. If the argument is well-constructed, the premises provide reasons in favour of the conclusion. In this example, the premises do seem to support the conclusion. At least they do, given the tacit assumption that you do not wish to get soaked.

This is the sort of thing that logicians are interested in. We shall say that an argument is any collection of premises, together with a conclusion.¹

In the example just given, we used individual sentences to express both of the argument’s premises, and we used a third sentence to express the argument’s conclusion. Many arguments are expressed in this way. But a single sentence can contain a complete argument. Consider:

I was wearing my sunglasses; so it must have been sunny.

¹ Because arguments are made of sentences, logicians are very concerned with the details of particular words and phrases appearing in sentences. Logic thus also has close connections with linguistics, particularly that subdiscipline of linguistics known as SEMANTICS, the theory of meaning.

This argument has one premise followed by a conclusion.

Many arguments start with premises, and end with a conclusion. But not all of them. The argument with which this section began might equally have been presented with the conclusion at the beginning, like so:

You should take an umbrella. After all, it is raining heavily. And if you do not take an umbrella, you will get soaked.

Equally, it might have been presented with the conclusion in the middle:

It is raining heavily. Accordingly, you should take an umbrella, given that if you do not take an umbrella, you will get soaked.

When approaching an argument, we want to know whether or not the conclusion follows from the premises. So the first thing to do is to separate out the conclusion from the premises. As a guideline, the following words are often used to indicate an argument's conclusion:

so, therefore, hence, thus, accordingly, consequently

And these expressions often indicate that we are dealing with a premise, rather than a conclusion

since, because, given that

But in analysing an argument, there is no substitute for a good nose.

Practice exercises

At the end of some sections, there are 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 memorising facts.

Highlight the phrase which expresses the conclusion of each of these arguments:

1. It is sunny. So I should take my sunglasses.
2. It must have been sunny. I did wear my sunglasses, after all.
3. No one but you has had their hands in the cookie-jar. And the scene of the crime is littered with cookie-crumbs. You're the culprit!
4. Miss Scarlett and Professor Plum were in the study at the time of the murder. And Reverend Green had the candlestick in the ballroom, and we know that there is no blood on his hands. Hence Colonel Mustard did it in the kitchen with the lead-piping. Recall, after all, that the gun had not been fired.

2 Valid arguments

In §1, we gave a very permissive account of what an argument is. To see just how permissive it is, consider the following:

There is a bassoon-playing dragon in Rundle Mall.
So: Salvador Dali was a poker player.

We have been given a premise and a conclusion. So we have an argument. Admittedly, it is a *terrible* argument. But it is still an argument.

2.1 Two ways that arguments can go wrong

It is worth pausing to ask what makes the argument so weak. In fact, there are two sources of weakness. First: the argument's (only) premise is obviously false. Rundle Mall has some interesting buskers, but not quite *that* interesting. Second: the conclusion does not follow from the premise of the argument. Even if there were a bassoon-playing dragon in Rundle Mall, we would not be able to draw any conclusion about Dali's predilection for poker.

What about the main argument discussed in §1? The premises of this argument might well be false. It might be sunny outside; or it might be that you can avoid getting soaked without taking an umbrella. But even if both premises were true, it does not necessarily show you that you should take an umbrella. Perhaps you enjoy walking in the rain, and you would like to get soaked. So, even if both premises were true, the conclusion might nonetheless be false. (If we were to add the formerly tacit assumption that you do not wish to get soaked as a further premise, then the premises taken together would provide support for the conclusion.)

The general point is as follows. For any argument, there are two ways that it might go wrong:

- › One or more of the premises might be false.
- › The conclusion might not follow from the premises – even if the premises were true, they would not support the conclusion.

To determine whether or not the premises of an argument are true is often a very important matter. But that is normally a task best left to experts in the field: as it might be, historians, scientists, or whomever. In our role as *logicians*, we are more concerned with arguments *in general*. So we are (usually) more concerned with the second way in which arguments can go wrong.

2.2 Reasons to believe

A good argument, in the logician's sense, links the premises to the conclusion. It turns the reasons you have for accepting the premises into reasons to accept its conclusion. But a good argument need not provide you with a reason to believe the conclusion. This can happen when you don't accept any of the premises in the first place. When the premises support the conclusion, that might just mean that they *would* be excellent reasons to accept the conclusion – *if only they were true!*

So: we are interested in whether or not a conclusion *follows from* some premises. Don't, though, say that the premises *infer* the conclusion. Entailment is a relation between premises and conclusions; inference is something we do. So if you want to mention inference when the conclusion follows from the premises, you could say that *one may infer* the conclusion from the premises.

But even this may be doubted. Often, when you believe the premises, a good argument provides you with a reason to believe the conclusion. In that case, it might be appropriate for you to infer the conclusion from the premises.

But sometimes a good argument shows that the premises support a conclusion you cannot accept. In that sort of situation, you might find that the argument gives you a better reason to abandon your belief in one of the premises than to infer the conclusion. A good argument shows there is some reason to believe its conclusion; it doesn't mean there aren't better reasons to reject its premises.

2.3 Conclusive arguments

As logicians, we want to be able to determine when the conclusion of an argument follows from the premises. One way to put this is as follows. We want to know whether, if all the premises were true, the conclusion would also have to be true. This motivates a definition:

An argument is **CONCLUSIVE** if and only if the truth of the premises guarantees the truth of the conclusion.

Consider another argument:

You are reading this book.

This is a logic book.

So: You are a logic student.

This is not a terrible argument. Both of the premises are true. And 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. This is not a conclusive argument.

The crucial thing about a conclusive argument is that it is impossible, in a very strict sense, for the premises to be true whilst the conclusion is false. Consider this example:

Oranges are either fruits or musical instruments.

Oranges are not fruits.

So: Oranges are musical instruments.

The conclusion of this argument is ridiculous. Nevertheless, it follows from the premises. *If* both premises were true, *then* the conclusion just has to be true. So the argument is conclusive.

Why is this argument conclusive? There are in fact two reasons we might point to:

- › The *structure* of the argument – the form of the premises and conclusion – ensures that the premises support the conclusion; and also
- › The premises cannot be true, and so of course the premises cannot be true while the conclusion is false.

The second reason applies to our argument because of the premise ‘Oranges are not fruits’. One might think it is *essential* to an orange that it be a fruit – many will say that *being a fruit* is part of what it is to be an orange. So ‘Oranges are not fruits’ is not only false: it is impossible for it to be true.

2.4 Conclusiveness for special reasons

There are plenty of arguments that are conclusive, but not for reasons relating to their structure. Take an example:

Juanita is a vixen.

So: Juanita is a fox.

It is impossible for the premise to be true and the conclusion false. So the argument is conclusive. But this is not due to the structure of the argument. Here is an inconclusive argument with seemingly the same form (it results from the first argument by replacing the word ‘fox’ with the word ‘cathedral’, but keeps the overall grammatical structure the same):

Juanita is a vixen.

So: Juanita is a cathedral.

This might suggest that the conclusiveness of the first argument *is* keyed to the meaning of the words ‘vixen’ and ‘fox’. But, whether or not that is right, it is not simply the *shape* of the argument that makes it conclusive. Equally, consider the argument:

The sculpture is green all over.

So: The sculpture is not red all over.

Again, because nothing can be both green all over and red all over, the truth of the premise would guarantee the truth of the conclusion. So the argument is conclusive. But here is an inconclusive argument with the same form:

The sculpture is green all over.

So: The sculpture is not shiny all over.

The argument is inconclusive, since it is possible to be green all over and shiny all over. (I might paint my nails with an elegant shiny green varnish.) Plausibly, the conclusiveness of this argument is keyed to the way that colours (or colour-words) interact. But, whether or not that is right, it is not simply the *shape* of the argument that makes it conclusive.

2.5 Validity

Logicians try to steer clear of controversial matters like whether there is a definition of an orange that makes it a fruit, or whether there is a ‘connection in meaning’ between being green and not being red. It is often difficult to figure such things out, and there may be widespread disagreement.

So logicians tend not to focus on conclusive arguments in general, but rather on those conclusive arguments which have a good structure or form. This is why the logic we are studying is sometimes called **FORMAL LOGIC**.

An argument is **VALID** if and only if it is conclusive due to its structure; otherwise it is **INVALID**.

The notion of the structure of a sentence, or an argument, is an intuitive one. I make the notion more precise in §4.1. Relying on our intuitive grasp of the notion for now, however, we can see this argument has the same form as the earlier argument about oranges, and is likewise valid:

EITHER Ogres are fearsome OR ogres are mythical.

IT IS NOT THE CASE THAT Ogres are fearsome.

So: Ogres are mythical.

This highlights that valid arguments do not need to have true premises (or even true conclusions). 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.
 So: Paris is in France.

The premises and conclusion of this argument are, as a matter of fact, all true. But the argument is invalid. If Paris were to declare independence from the rest of France, then the conclusion would be false, even though both of the premises would remain true. Thus, it is *possible* for the premises of this argument to be true and the conclusion false. The argument is therefore inconclusive, and hence invalid.

The important thing to remember is that validity is not about the actual truth or falsity of the sentences in the argument. It is about whether the structure of the argument ensures that the premises support the conclusion. Nonetheless, we shall say that an argument is **SOUND** if and only if it is both valid and all of its premises are true.

2.6 Inductive arguments

Many good arguments are inconclusive and invalid. Consider this one:

In January 1997, it rained in London.
 In January 1998, it rained in London.
 In January 1999, it rained in London.
 In January 2000, it rained in London.
 So: It rains every January in London.

This argument generalises from observations about several cases to a conclusion about all cases. Such arguments are called **INDUCTIVE** arguments. The argument could be made stronger by adding additional premises before drawing the conclusion: In January 2001, it rained in London; In January 2002.... But, however many premises of this form we add, the argument will remain inconclusive. Even if it has rained in London in every January thus far, it remains *possible* that London will stay dry next January.

The point of all this is that inductive arguments – even good inductive arguments – are not (deductively) valid. They are not *watertight*. Unlikely though it might be, it is *possible* for their conclusion to be false, even when all of their premises are true. In this book, we shall set aside (entirely) the question of what makes for a good inductive argument. Our interest is simply in sorting the (deductively) valid arguments from the invalid ones.

Practice exercises

A. Which of the following arguments are valid? Which are invalid but conclusive? Which are inconclusive? Comment on any difficulties or points of interest.

1. Socrates is a man.
 2. All men are carrots.
- So: Therefore, Socrates is a carrot.

1. Abe Lincoln was either born in Illinois or he was once president.
2. Abe Lincoln was never president.

So: Abe Lincoln was born in Illinois.

1. Abe Lincoln was president.

So: Abe Lincoln was a citizen of the United States.

1. If I pull the trigger, Abe Lincoln will die.
2. I do not pull the trigger.

So: Abe Lincoln will not die.

1. Abe Lincoln was either from France or from Luxemborg.
2. Abe Lincoln was not from Luxemborg.

So: Abe Lincoln was from France.

1. If the world were to end today, then I would not need to get up tomorrow morning.
2. I will need to get up tomorrow morning.

So: The world will not end today.

1. Joe is now 19 years old.
2. Joe is now 87 years old.

So: Bob is now 20 years old.

B. Could there be:

1. A valid argument that has one false premise and one true premise?
2. A valid argument that has only false premises?
3. A valid argument with only false premises and a false conclusion?
4. A sound argument with a false conclusion?
5. An invalid argument that can be made valid by the addition of a new premise?
6. A valid argument that can be made invalid by the addition of a new premise?

In each case: if so, give an example; if not, explain why not.

3 Other logical notions

In §2, we introduced the idea of a valid argument. We will want to introduce some more ideas that are important in logic.

3.1 Truth values

As we said in §1, arguments consist of premises and a conclusion. Note that many kinds of English sentence cannot be used to express premises or conclusions of arguments. For example:

- › **Questions**, e.g., ‘are you feeling sleepy?’
- › **Imperatives**, e.g., ‘Wake up!’
- › **Exclamations**, e.g., ‘Ouch!’

The common feature of these three kinds of sentence is that they cannot be used to make *assertions*: they cannot be true or false. It does not even make sense to ask whether a *question* is true (it only makes sense to ask whether the *answer* to a question is true).

The general point is that, the premises and conclusion of an argument must be capable of having a TRUTH VALUE. And the two truth values that concern us are just True and False.

3.2 Consistency

Consider these two sentences:

- B₁. Jane’s only brother is shorter than her.
- B₂. Jane’s only brother is taller than her.

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 is impossible that both sentences are true together. These sentences are inconsistent with each other. And this motivates the following definition:

Sentences are JOINTLY CONSISTENT if and only if it is possible for them all to be true together.

Conversely, B₁ and B₂ are JOINTLY INCONSISTENT.

We can ask about the consistency of any number of sentences. For example, consider the following four 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 entail 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 sentences G₁–G₄ are jointly inconsistent. They cannot all be true together. (Note that the sentences G₁, G₃ and G₄ are jointly inconsistent. But if sentences are already jointly inconsistent, adding an extra sentence to the mix will not make them consistent!)

There is an interesting connection between consistency and conclusive arguments. A conclusive argument is one where the premises guarantee the truth of the conclusion. So it is an argument where if the premises are true, the conclusion must be true. So the premises cannot be jointly consistent with the claim that the conclusion is false. Since the argument ‘Dogs and cats are animals, so dogs are animals’ is conclusive, that shows that the sentences ‘Dogs and cats are animals’ and ‘Dogs are **not** animals’ are jointly inconsistent. If an argument is conclusive, the premises of the argument taken together with the denial of the conclusion will be jointly inconsistent.

We just linked consistency to conclusive arguments. There is an analogous notion linked to valid arguments:

Sentences are JOINTLY FORMALLY CONSISTENT if and only if, considering only their structure, they can all be true together.

If some sentences are jointly consistent, they are also jointly formally consistent. But some formally consistent sentences are jointly inconsistent. Conclusive but invalid argument will give us an example. Since ‘The sculpture is green all over, so the sculpture is not red all over’ is conclusive, the sentences ‘The sculpture is green’ and ‘The sculpture is red’ are jointly inconsistent. But those sentences are formally consistent – their inconsistency depends crucially on the meanings of the words ‘red’ and ‘green’, not on the overall structure of the sentences in which those words appear. (Again, more on the notion of structure invoked here in §4.1.)

3.3 Necessity and contingency

In assessing whether an argument is conclusive, we care about what would be true *if* the premises were true. But some sentences just *must* be true. Consider these sentences:

1. It is raining.
2. If it is raining, water is precipitating from the sky.
3. Either it is raining here, or it is not.
4. It is both raining here and not raining here.

In order to know if sentence **1** is true, you would need to look outside or check the weather channel. It might be true; it might be false.

Sentence **2** is different. You do not need to look outside to know that it says something true. Regardless of what the weather is like, if it is raining, water is precipitating – that is just what rain *is*, meteorologically speaking. That is a NECESSARY TRUTH. Here, a necessary connection in meaning between ‘rain’ and ‘precipitation’ makes what the sentence says true in every circumstance.

Sentence **3** is also a necessary truth. Unlike sentence **2**, however, it is the structure of the sentence which makes it necessary. No matter what ‘raining here’ means, ‘Either it is raining here or it is not raining here’ will be true. The structure ‘EITHER IT IS ... OR IT IS NOT ...’, where both gaps (...) are filled by the same phrase, must be a true sentence.

Equally, you do not need to check the weather, or even the meaning of words, to determine whether or not sentence **4** is true. It must be false, simply as a matter of structure. It might be raining here and not raining across town; it might be raining now but stop raining even as you finish this sentence; but it is impossible for it to be both raining and not raining in the same place and at the same time. So, whatever the world is like, it is not both raining here and not raining here. It is a NECESSARY FALSEHOOD.

A sentence which is capable of being true or false, but which says something which is neither necessarily true nor necessarily false, is CONTINGENT.

If a sentence says something which is sometimes true and sometimes false, it will definitely be contingent. But something might *always* be true and still be contingent. For instance, it seems plausible that whenever there have been people, some of them habitually arrive late. ‘Some people are habitually late’ is always true. But it is contingent, it seems: human nature could have been more punctual. If so, the sentence would have been false. But if something is really necessary, it will always be true.¹

If some sentences contain amongst themselves a necessary falsehood, those sentences are jointly inconsistent. At least one of them cannot be true, so they cannot all be true together. Accordingly, if an argument has a premise that is a necessary falsehood, or its

¹ Here’s an interesting example to consider. It seems that, whenever anyone says the sentence ‘I am here now’, they say something true. That sentence is, whenever it is uttered, *truly uttered*. But does it say something necessary or contingent?

conclusion is a necessary truth, or both, then the argument is conclusive – its premises and the denial of its conclusion will be jointly inconsistent.

Practice exercises

A. For each of the following: Is it necessarily true, necessarily false, or contingent?

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.

B. Look back at the sentences G_1 – G_4 in this section (about giraffes, gorillas and martians in the wild animal park), and consider each of the following:

1. G_2 , G_3 , and G_4
2. G_1 , G_3 , and G_4
3. G_1 , G_2 , and G_4
4. G_1 , G_2 , and G_3

Which are jointly consistent? Which are jointly inconsistent?

C. Could there be:

1. A conclusive argument, the conclusion of which is necessarily false?
2. An inconclusive argument, the conclusion of which is necessarily true?
3. Jointly consistent sentences, one of which is necessarily false?
4. Jointly inconsistent sentences, one of which is necessarily true?

In each case: if so, give an example; if not, explain why not.

Chapter 2

Sentential logic

4 First steps to symbolisation

4.1 Argument Structure

Consider this argument:

It is raining outside.

If it is raining outside, then Jenny is miserable.

So: Jenny is miserable.

and another argument:

Jenny is an anarcho-syndicalist.

If Jenny is an anarcho-syndicalist, then Dipan is an avid reader of Tolstoy.

So: Dipan is an avid reader of Tolstoy.

Both arguments are valid, and there is a straightforward sense in which we can say that they share a common structure. We might express the structure thus:

A

If A, then C

So: C

This is an excellent argument *structure*. Surely any argument with this *structure* will be valid. And this is not the only good argument structure. Consider an argument like:

Jenny is either happy or sad.

Jenny is not happy.

So: Jenny is sad.

Again, this is a valid argument. The structure here is something like:

A or B

not: A

So: B

A superb structure! And here is a final example:

It's not the case that Jim both studied hard and acted in lots of plays.
 Jim studied hard
 So: Jim did not act in lots of plays.

This valid argument has a structure which we might represent thus:

not both: A and B
 A
 So: not: B

The examples illustrate the idea of validity – conclusiveness in virtue of structure. The validity of the arguments just considered has nothing very much to do with the meanings of English expressions like 'Jenny is miserable', 'Dipan is an avid reader of Tolstoy', or 'Jim acted in lots of plays'. If it has to do with meanings at all, it is with the meanings of phrases like 'and', 'or', 'not,' and 'if ..., then ...'

In this chapter, we are going to develop a formal language which allows us to symbolise many arguments in such a way as to show that they are valid, conclusive in virtue of their form. That language will be called *Sentential*, and the study of that language and its features is called *sentential logic*. (It also has other names: see Appendix A.)

Not every conclusive argument is valid. Arguments with necessary conclusions or impossible premises will definitely be conclusive, but need not be valid. And arguments whose conclusiveness does turn on the meanings of words might not be valid either. An important moral can be stated as follows. *At best, logic will help us to understand arguments that are valid, or conclusive due to their form.* Luckily, the great majority of conclusive arguments either are valid, or can be transformed into valid arguments by making some assumption they implicitly rely on into an explicit premise.

For example, recall this conclusive argument:

Juanita is a vixen.
 So: Juanita is a fox.

It can be transformed into a valid argument by making the implicit assumption that *all vixens are foxes* into an appropriate premise:

Juanita is a vixen.
 If Juanita is a vixen, then Juanita is a fox.
 So: Juanita is a fox.

This argument is valid – it shares a structure with the first two arguments we considered in this chapter.

4.2 Identifying Argument Structure

How can we identify the logical structure of an argument? The logical structure of an argument is the form one arrives at by ‘abstracting away’ the details of the words in an argument, *except for a special set of words*. In the examples above, we removed all details of the arguments except for the special words ‘and’, ‘or’, ‘not’ and ‘if ..., then ...’, and replaced the other phrases in the argument by placeholder letters ‘A’, ‘B’, etc. (We were careful to replace the same phrase by the same placeholder)

This raises another question. What makes those words special? Logicians tend to take a pragmatic attitude to this question. They say: *nothing!* If you had chosen different words, you would have come up with a different structure. But there are practical reasons why the words I’ve already highlighted are useful ones – primarily because the formal language that is able to model arguments with this structure is powerful, useful, and nicely behaved.

But we will already see later in this book that one could focus on different aspects of the argument as structural. The language we develop in chapter 4 is one that also takes names like ‘Juanita’ and predicates like ‘is a vixen’ or ‘is a fox’ to be structurally important. And there are still other formal logical languages, unfortunately beyond the scope of this book, which take still other words or grammatical categories as structural: logics which take modal words like ‘necessarily’ or ‘possibly’ as structural, and logics which take temporal adverbs like ‘always’ and ‘now’ as structural.

But for now, put those tantalising observations out of your mind! We’ll focus to begin with on the language `Sentential`, and on how we can use it to model arguments in English featuring some particular structural words, including those we’ve just highlighted above.

4.3 Atomic sentences

A `SUBSENTENCE` of a sentence is a phrase occurring in a sentence which is itself a grammatical sentence. I started isolating the form of an argument, in §4.1, by replacing subsentences of sentences with individual placeholder letters. Thus in the first example of this section, ‘it is raining outside’ is a subsentence of ‘If it is raining outside, then Jenny is miserable’, and we replaced this subsentence with ‘A’.

Our artificial language, `Sentential`, pursues this idea absolutely ruthlessly. We start with some *atomic sentences*. These will be the basic building blocks out of which more complex sentences are built. We will use upper case italic letters for atomic sentences of `Sentential`. There are only twenty-six letters of the alphabet, but there is no limit to the number of atomic sentences that we might want to consider. By adding subscripts to letters, we obtain new atomic sentences. So, here are five different atomic sentences of `Sentential`:

$$A, P, P_1, P_2, A_{234}.$$

We shall use atomic sentence to represent, or symbolise, certain English sentences. To do this, we provide a `SYMBOLISATION KEY`, such as the following:

A: It is raining outside.
 C: Jenny is miserable.

In doing this, we are not fixing this symbolisation *once and for all*. We are just saying that, for the time being, we shall use the atomic sentence of Sentential, 'A', to symbolise the English sentence 'It is raining outside', and the atomic sentence of Sentential, 'C', to symbolise the English sentence 'Jenny is miserable. Later, when we are dealing with different sentences or different arguments, we can provide a new symbolisation key; as it might be:

A: Jenny is an anarcho-syndicalist.
 C: Dipan is an avid reader of Tolstoy.

But it is important to understand that whatever structure an English sentence might have is lost when it is symbolised by an atomic sentence of Sentential. From the point of view of Sentential, an atomic sentence is just a letter. It can be used to build more complex sentences, but it cannot be taken apart. So we cannot use Sentential to symbolise arguments whose validity depends on structure within atomic sentences.

4.4 Structure and Ambiguity

In later sections, we will see the utility of symbolising arguments to evaluate their validity. But even at this early stage, there is value in symbolising sentences, because doing so can clear up *ambiguity*.

Consider the following imagined regulation: 'Small children must be quiet and seated or carried at all times'. This sentence has an ambiguous structure, by which I mean: there are two different grammatical structures that it might be thought to have. Symbolising brings this out. Let us use the following key:

Q: Small children must be quiet.
 S: Small children must be seated.
 C: Small children must be carried.

Leaving the temporal adverbial phrase 'at all times' aside, here are two symbolisations of our target sentence:

Q AND EITHER S OR C
 EITHER Q AND S, OR C

This ambiguity matters. If a child is carried, but is not quiet, the parent is violating the regulation with the first structure, but in compliance with the regulation with the second structure. As we will soon see, Sentential has resources to ensure that no ambiguity is present in any symbolisation of a given target sentence, which helps us make clear what we really might have meant by that sentence.

5 Connectives

In the previous section, we considered symbolising fairly basic English sentences with atomic sentences of `Sentential`. This leaves us wanting to deal with the English expressions ‘and’, ‘or’, ‘not’, and so forth. These are *connectives* – they can be used to form new sentences out of old ones. And in `Sentential`, we shall make use of logical connectives to build complex sentences from atomic components. There are five logical connectives in `Sentential`. This table summarises them, and they are explained throughout this section.

symbol	what it is called	rough meaning
\neg	negation	‘It is not the case that ...’
\wedge	conjunction	‘Both ... and ...’
\vee	disjunction	‘Either ... or ...’
\rightarrow	conditional	‘If ... then ...’
\leftrightarrow	biconditional	‘... if and only if ...’

5.1 Negation

Consider how we might symbolise these sentences:

1. Mary is in Barcelona.
2. It is not the case that Mary is in Barcelona.
3. Mary is not in Barcelona.

In order to symbolise sentence **1**, we will need an atomic sentence. We might offer this symbolisation key:

B : Mary is in Barcelona.

Since sentence **2** is obviously related to the sentence **1**, we shall not want to symbolise it with a completely different sentence. Roughly, sentence **2** means something like ‘It is not the case that B ’. In order to symbolise this, we need a symbol for negation. We will use ‘ \neg ’. Now we can symbolise sentence **2** with ‘ $\neg B$ ’.

Sentence 3 also contains the word ‘not’. And it is obviously equivalent to sentence 2. As such, we can also symbolise it with ‘ $\neg B$ ’.

It is much more common in English to see negation appear as it does in 3, with a ‘not’ found somewhere within the sentence, than the more formal form in 2. The form in 2 has the benefit that the sentence ‘Mary is in Barcelona’ itself appears as a subsentence of 1 – so we can see how negation forms a new sentence by literally adding words to the old sentence. This is not so direct in 3. But since 2 and 3 are near enough synonymous in meaning, we can treat them both as negations of 1. (This issue is a bit tricky – see the discussion of examples 7 and 8.)

If a sentence can be paraphrased in English as ‘It is not the case that ...’, then it may be symbolised by $\neg A$, where ‘ A ’ is the Sentential sentence symbolising the sentence occurring at ‘...’¹

It will help to offer a few more examples:

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

Let us use the following representation key:

R : The widget is replaceable.

Sentence 4 can now be symbolised by ‘ R ’. Moving on to sentence 5: saying the widget is irreplaceable means that it is not the case that the widget is replaceable. So even though sentence 5 does not contain the word ‘not’, we shall symbolise it as follows: ‘ $\neg R$ ’.

Sentence 6 can be paraphrased as ‘It is not the case that the widget is irreplaceable.’ Which can again be paraphrased as ‘It is not the case that it is not the case that the widget is replaceable.’ So we might symbolise this English sentence with the Sentential sentence ‘ $\neg\neg R$ ’.

But some care is needed when handling negations. Consider:

7. Jane is happy.
8. Jane is unhappy.

If we let the Sentential-sentence ‘ H ’ symbolise ‘Jane is happy’, then we can symbolise sentence 7 as ‘ H ’. However, it would be a mistake to symbolise sentence 8 with ‘ $\neg H$ ’. If Jane is unhappy, then she is not happy; but sentence 8 does not mean the same thing as ‘It is not the case that Jane is happy’. Jane might be neither happy nor unhappy; she might be in a state of blank indifference. In order to symbolise sentence 8, then, we would need a new atomic sentence of Sentential.

5.2 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 atomic sentences of Sentential to symbolise sentences 9 and 10; perhaps

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

Sentence 9 can now be symbolised as ' A ', and sentence 10 can be symbolised as ' B '. Sentence 11 roughly says 'A and B'. We need another symbol, to deal with 'and'. We will use ' \wedge '. Thus we will symbolise it as ' $(A \wedge B)$ '. This connective is called CONJUNCTION. We also say that ' A ' and ' B ' are the two CONJUNCTS of the conjunction ' $(A \wedge B)$ '.

Notice that we make no attempt to symbolise the word 'also' in sentence 11. Words like 'both' and 'also' function to draw our attention to the fact that two things are being conjoined. Maybe they affect the emphasis of a sentence. But we will not (and cannot) symbolise such things in Sentential.

Some more examples will bring out this point:

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

Sentence 12 is obviously a conjunction. The sentence says two things (about Barbara). In English, it is permissible to refer to Barbara only once. It *might* be tempting to think that we need to symbolise sentence 12 with something along the lines of ' B and energetic'. This would be a mistake. Once we symbolise part of a sentence as ' B ', any further structure is lost. ' B ' is an atomic sentence of Sentential. Conversely, 'energetic' is not an English sentence at all. What we are aiming for is something like ' B and Barbara is energetic'. So we need to add another sentence letter to the symbolisation key. Let ' E ' symbolise 'Barbara is energetic'. Now the entire sentence can be symbolised as ' $(B \wedge E)$ '.

Sentence 13 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. The sentence can be paraphrased as 'Barbara is athletic, and Adam is athletic'. We can symbolise this in Sentential as ' $(B \wedge A)$ ', using the same symbolisation key that we have been using.

Sentence 14 is slightly more complicated. The word 'although' sets up a contrast between the first part of the sentence and the second part. Nevertheless, the sentence tells us 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'. And now we can symbolise this with the Sentential sentence ' $(E \wedge \neg B)$ '. Note that we have lost all sorts of nuance in this symbolisation. There is a distinct difference in tone between sentence 14 and 'Both Barbara is energetic and it is not the case that Barbara is athletic'. Sentential does not (and cannot) preserve these nuances.

Sentence 15 raises similar issues. There is a contrastive structure, but this is not something that Sentential can deal with. So we can paraphrase the sentence as 'Both Adam is athletic, and Barbara is more athletic than Adam'. (Notice that we once again replace the pronoun 'him' with 'Adam'.) How should we deal with the second conjunct? We already have the sentence letter 'A', which is being used to symbolise 'Adam is athletic', and the sentence 'B' which is being used to symbolise 'Barbara is athletic'; but neither of these concerns their relative athleticism. So, to symbolise the entire sentence, we need a new sentence letter. Let the Sentential sentence 'R' symbolise the English sentence 'Barbara is more athletic than Adam'. Now we can symbolise sentence 15 by ' $(A \wedge R)$ '.

A sentence can be symbolised as $(A \wedge B)$ if it can be paraphrased in English as 'Both ..., and ...', or as '..., but ...', or as 'although ..., ...'

You might be wondering why I am putting brackets around the conjunctions. The reason for this is to avoid potential ambiguity. This can be brought out by considering how negation might interact with conjunction. Consider:

16. It's not the case that you will get both soup and salad.
17. You will not get soup but you will get salad.

Sentence 16 can be paraphrased as 'It is not the case that: both you will get soup and you will get salad'. Using this symbolisation key:

- S_1 : You get soup.
 S_2 : You get salad.

We would symbolise 'both you will get soup and you will get salad' as ' $(S_1 \wedge S_2)$ '. To symbolise sentence 16, then, we simply negate the whole sentence, thus: ' $\neg(S_1 \wedge S_2)$ '.

Sentence 17 is a conjunction: you *will not* get soup, and you *will* get salad. 'You will not get soup' is symbolised by ' $\neg S_1$ '. So to symbolise sentence 17 itself, we offer ' $(\neg S_1 \wedge S_2)$ '.

These English sentences are very different, and their symbolisations differ accordingly. In one of them, the entire conjunction is negated. In the other, just one conjunct is negated. Brackets help us to avoid ambiguity by clear distinguishing these two cases.

Again, English is subject to this sort of ambiguity. Suppose instead of 16, we'd just said

18. It is not the case that you will get soup and salad.

The sentence 18 is ambiguous; one of its readings says the same thing as 16, the other says the same thing as 17. Brackets enable Sentential to avoid precisely this ambiguity.

The introduction of brackets prompts us to define some other useful concepts. If a sentence of Sentential has the form $(\mathcal{A} \wedge \mathcal{B})$, we say that its *main connective* is conjunction – even if other connectives occur within \mathcal{A} or \mathcal{B} . Likewise, if the sentence has the form $\neg\mathcal{A}$, its main connective is negation. We define the *scope* of an occurrence of a connective in a sentence as the subsentence which has that connective as its main connective. So the scope of ‘ \neg ’ in ‘ $\neg(S_1 \wedge S_2)$ ’ is the whole sentence (because negation is the main connective), while the scope of ‘ \neg ’ in ‘ $(\neg S_1 \wedge S_2)$ ’ is just the subsentence ‘ $\neg S_1$ ’ – the main connective of the whole sentence is conjunction. Brackets help us keep track of things the scope of our connectives, and which connective in a sentence is the main connective. I say more about the notions of a main connective and the scope of a connective in chapter 6.

5.3 Disjunction

Consider these sentences:

- 19. Either Denison will play golf with me, or he will watch movies.
- 20. Either Denison or Ellery will play golf with me.

For these sentences we can use this symbolisation key:

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

However, we shall again need to introduce a new symbol. Sentence 19 is symbolised by ‘ $(D \vee M)$ ’. The connective is called DISJUNCTION. We also say that ‘ D ’ and ‘ M ’ are the DISJUNCTS of the disjunction ‘ $(D \vee M)$ ’.

Sentence 20 is only slightly more complicated. There are two subjects, but the English sentence only gives the verb once. However, we can paraphrase sentence 20 as ‘Either Denison will play golf with me, or Ellery will play golf with me’. Now we can obviously symbolise it by ‘ $(D \vee E)$ ’ again.

A sentence can be symbolised as $(\mathcal{A} \vee \mathcal{B})$ if it can be paraphrased in English as ‘Either ..., or ...’. 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 20, above. I might play golf with Denison, with Ellery, or with both Denison and Ellery. Sentence 20 merely says that I will play with *at least* one of them. This is called an INCLUSIVE OR. The Sentential symbol ‘ \vee ’ always symbolises an inclusive ‘or’.

It might help to see negation interact with disjunction. Consider:

21. Either you will not have soup, or you will not have salad.
22. You will have neither soup nor salad.
23. You get either soup or salad, but not both.

Using the same symbolisation key as before, sentence 21 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’. To symbolise this in Sentential, we need both disjunction and negation. ‘It is not the case that you get soup’ is symbolised by ‘ $\neg S_1$ ’. ‘It is not the case that you get salad’ is symbolised by ‘ $\neg S_2$ ’. So sentence 21 itself is symbolised by ‘ $(\neg S_1 \vee \neg S_2)$ ’.

Sentence 22 also requires negation. It can be paraphrased as, ‘*It is not the case that* either you get soup or you get salad’. Since this negates the entire disjunction, we symbolise sentence 22 with ‘ $\neg(S_1 \vee S_2)$ ’.

Sentence 23 is an exclusive ‘or’. We can break the sentence into two parts. The first part says that you get one or the other. We symbolise 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 symbolise this with ‘ $\neg(S_1 \wedge S_2)$ ’. Now we just need to put the two parts together. As we saw above, ‘but’ can usually be symbolised with ‘ \wedge ’. Sentence 23 can thus be symbolised as ‘ $((S_1 \vee S_2) \wedge \neg(S_1 \wedge S_2))$ ’. This last example shows something important. Although the Sentential symbol ‘ \vee ’ always symbolises inclusive ‘or’, we can symbolise an exclusive ‘or’ in Sentential. We just have to use a few of our other symbols as well.

5.4 Conditional

Consider these sentences:

24. If Jean is in Paris, then Jean is in France.
25. Jean is in France only if Jean is in Paris.

Let’s use the following symbolisation key:

- P : Jean is in Paris.
 F : Jean is in France.

Sentence 24 is roughly of this form: ‘if P , then F ’. We will use the symbol ‘ \rightarrow ’ to symbolise this ‘if ..., then ...’ structure. So we symbolise sentence 24 by ‘ $(P \rightarrow F)$ ’. The connective is

called THE CONDITIONAL. Here, ‘ P ’ is called the ANTECEDENT of the conditional ‘ $(P \rightarrow F)$ ’, and ‘ F ’ is called the CONSEQUENT.

Sentence 25 is also a conditional. Since the word ‘if’ appears in the second half of the sentence, it might be tempting to symbolise this in the same way as sentence 24. That would be a mistake. My knowledge of geography tells me that sentence 24 is unproblematically true: there is no way for Jean to be in Paris that doesn’t involve Jean being in France. But sentence 25 is not so straightforward: were Jean in Dijon, Marseilles, or Toulouse, Jean would be in France without being in Paris, thereby rendering sentence 25 false. Since geography alone dictates the truth of sentence 24, whereas travel plans (say) are needed to know the truth of sentence 25, they must mean different things.

In fact, sentence 25 can be paraphrased as ‘If Jean is in France, then Jean is in Paris’. So we can symbolise it by ‘ $(F \rightarrow P)$ ’.

A sentence can be symbolised as $\mathcal{A} \rightarrow \mathcal{B}$ if it can be paraphrased in English as ‘If A, then B’ or ‘A only if B’.

In fact, many English expressions can be represented using the conditional. Consider:

26. For Jean to be in Paris, it is necessary that Jean be in France.
27. It is a necessary condition on Jean’s being in Paris that she be in France.
28. For Jean to be in France, it is sufficient that Jean be in Paris.
29. It is a sufficient condition on Jean’s being in France that she be in Paris.

If we think really hard, all four of these sentences mean the same as ‘If Jean is in Paris, then Jean is in France’. So they can all be symbolised by ‘ $P \rightarrow F$ ’.

It is important to bear in mind that the connective ‘ \rightarrow ’ tells us only that, if the antecedent is true, then the consequent is true. It says nothing about a *causal* connection between two events (for example). In fact, we seem to lose a huge amount when we use ‘ \rightarrow ’ to symbolise English conditionals. We shall return to this in §§9.3 and 11.5.

5.5 Biconditional

Consider these sentences:

30. Shergar is a horse only if it he is a mammal.
31. Shergar is a horse if he is a mammal.
32. Shergar is a horse if and only if he is a mammal.

We shall use the following symbolisation key:

- H : Shergar is a horse.
 M : Shergar is a mammal.

Sentence 30, for reasons discussed above, can be symbolised by ' $H \rightarrow M$ '.

Sentence 31 is importantly different. It can be paraphrased as, 'If Shergar is a mammal then Shergar is a horse'. So it can be symbolised by ' $M \rightarrow H$ '.

Sentence 32 says something stronger than either 30 or 31. It can be paraphrased as 'Shergar is a horse if he is a mammal, and Shergar is a horse only if Shergar is a mammal'. This is just the conjunction of sentences 30 and 31. So we can symbolise it as ' $(H \rightarrow M) \wedge (M \rightarrow H)$ '. We call this a BICONDITIONAL, because it entails the conditional in both directions.

We could treat every biconditional this way. So, just as we do not need a new Sentential symbol to deal with exclusive 'or', we do not really need a new Sentential symbol to deal with biconditionals. However, we will use ' \leftrightarrow ' to symbolise the biconditional. So we can symbolise sentence 32 with the Sentential sentence ' $H \leftrightarrow M$ '.

The expression 'if and only if' occurs a lot in philosophy and logic. For brevity, we can abbreviate it with the snappier word 'iff'. I shall follow this practice. So 'if' with only *one* 'f' is the English conditional. But 'iff' with *two* 'f's is the English biconditional. Armed with this we can say:

A sentence can be symbolised as $\mathcal{A} \leftrightarrow \mathcal{B}$ if it can be paraphrased in English as 'A iff B'; that is, as 'A if and only if B'.

Other expressions in English which can be used to mean 'iff' include 'exactly if' and 'exactly when', or even 'just in case'. So if we say 'You run out of time exactly when the buzzer sounds', we mean: 'if the buzzer sounds, then you are out of time; and also if you are out of time, then the buzzer sounds'.

A word of caution. Ordinary speakers of English often use 'if ..., then ...' when they really mean to use something more like '... if and only if ...'. Perhaps your parents told you, when you were a child: 'if you don't eat your vegetables, you won't get any dessert'. Suppose you ate your vegetables, but that your parents refused to give you any dessert, on the grounds that they were only committed to the *conditional* (roughly 'if you get dessert, then you will have eaten your vegetables'), rather than the biconditional (roughly, 'you get dessert iff you eat your vegetables'). Well, a tantrum would rightly ensue. So, be aware of this when interpreting people; but in your own writing, make sure you use the biconditional iff you mean to.

5.6 Unless

We have now introduced all of the connectives of Sentential. We can use them together to symbolise many kinds of sentences. But a typically nasty case is when we use the English-language connective 'unless':

33. Unless you wear a jacket, you will catch cold.
34. You will catch cold unless you wear a jacket.

These two sentences are clearly equivalent. To symbolise them, we shall use the symbolisation key:

J: You will wear a jacket.

D: You will catch a cold.

Both sentences mean that if you do not wear a jacket, then you will catch cold. With this in mind, we might symbolise them as ' $\neg J \rightarrow D$ '.

Equally, both sentences mean that if you do not catch a cold, then you must have worn a jacket. With this in mind, we might symbolise them as ' $\neg D \rightarrow J$ '.

Equally, both sentences mean that either you will wear a jacket or you will catch a cold. With this in mind, we might symbolise them as ' $J \vee D$ '.

All three are correct symbolisations. Indeed, in chapter 3 we shall see that all three symbolisations are equivalent in Sentential.

If a sentence can be paraphrased as 'Unless A, B' then it can be symbolised as ' $A \vee B$ '.

Again, though, there is a little complication. 'Unless' can be symbolised as a conditional; but as I said above, people often use the conditional (on its own) when they mean to use the biconditional. Equally, 'unless' can be symbolised as a disjunction; but there are two kinds of disjunction (exclusive and inclusive). So it will not surprise you to discover that ordinary speakers of English often use 'unless' to mean something more like the biconditional, or like exclusive disjunction. Suppose I say: 'I shall go running unless it rains'. I probably mean something like 'I shall go running iff it does not rain' (i.e., the biconditional), or 'either I shall go running or it will rain, but not both' (i.e., exclusive disjunction). Again: be aware of this when interpreting what other people have said, but be precise in your writing, unless you want to be deliberately ambiguous.

Practice exercises

A. Using the symbolisation key given, symbolise each English sentence in Sentential.

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.

B. Using the symbolisation key given, symbolise each English sentence in Sentential.

- 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. Since the cook did it, the butler did not.
12. Of course the Duchess is lying!

C. Using the symbolisation key given, symbolise each English sentence in Sentential.

- E_1 : Ava is an electrician.
 E_2 : Harrison is an electrician.
 F_1 : Ava is a firefighter.
 F_2 : Harrison is a firefighter.
 S_1 : Ava is satisfied with her career.
 S_2 : 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, 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.

D. Give a symbolisation key and symbolise the following English sentences in Sentential.

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.

E. Give a symbolisation key and symbolise the following English sentences in Sentential.

1. If there is food to be found in the pridelands, then Rafiki will talk about squashed bananas.
2. Rafiki will talk about squashed bananas unless Simba is alive.
3. Rafiki will either talk about squashed bananas or he won't, but there is food to be found in the pridelands regardless.
4. Scar will remain as king if and only if there is food to be found in the pridelands.
5. If Simba is alive, then Scar will not remain as king.

F. For each argument, write a symbolisation key and symbolise all of the sentences of the argument in Sentential.

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.

G. We symbolised an exclusive 'or' using ' \vee ', ' \wedge ', and ' \neg '. How could you symbolise an exclusive 'or' using only two connectives? Is there any way to symbolise an exclusive 'or' using only one connective?

6 Sentences of Sentential

The sentence ‘either apples are red, or berries are blue’ is a sentence of English, and the sentence ‘ $(A \vee B)$ ’ is a sentence of Sentential. Although we can identify sentences of English when we encounter them, we do not have a formal definition of ‘sentence of English’. But in this chapter, we shall offer a complete *definition* of what counts as a sentence of Sentential. This is one respect in which a formal language like Sentential is more precise than a natural language like English. Of course, Sentential was designed to be much simpler than English.

6.1 Expressions

We have seen that there are three kinds of symbols in Sentential:

Atomic sentences with subscripts, as needed	A, B, C, \dots, Z $A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \dots$
Connectives	$\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
Brackets	$(,)$

We define an **EXPRESSION** of Sentential as any string of symbols of Sentential. Take any of the symbols of Sentential and write them down, in any order, and you have an expression of Sentential. Expressions are sometimes called ‘strings’, because they are just a string of symbols from the approved list above.

6.2 Sentences

We want to know when an expression of Sentential amounts to a *sentence*. Many expressions of Sentential will be uninterpretable nonsense. ‘ $(A_{17}JQF\neg K))\wedge()$ ’ is a perfectly good expression, but doesn’t look likely to end up a correctly formed sentence of our language. Accordingly, we need to clarify the grammatical rules of Sentential. We’ve already seen some of those rules when we introduced the atomic sentences and connectives. But I will now make them all explicit.

Obviously, individual atomic sentences like ' A ' and ' G_{13} ' should count as sentences. We can form further sentences out of these by using the various connectives. Using negation, we can get ' $\neg A$ ' and ' $\neg G_{13}$ '. Using conjunction, we can get ' $(A \wedge G_{13})$ ', ' $(G_{13} \wedge A)$ ', ' $(A \wedge A)$ ', and ' $(G_{13} \wedge G_{13})$ '. We could also apply negation repeatedly to get sentences like ' $\neg\neg A$ ' or apply negation along with conjunction to get sentences like ' $\neg(A \wedge G_{13})$ ' and ' $\neg(G_{13} \wedge \neg G_{13})$ '. 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 sentences one by one.

Instead, we will describe the process by which sentences can be *constructed*. Consider negation: Given any sentence \mathcal{A} of Sentential, $\neg\mathcal{A}$ is a sentence of Sentential. (Why the funny fonts? I return to this in §7.)

We can say similar things for each of the other connectives. For instance, if \mathcal{A} and \mathcal{B} are sentences of Sentential, then $(\mathcal{A} \wedge \mathcal{B})$ is a sentence of Sentential. Providing clauses like this for all of the connectives, we arrive at the following formal definition for a SENTENCE of Sentential

1. Every atomic sentence is a sentence.
2. If \mathcal{A} is a sentence, then $\neg\mathcal{A}$ is a sentence.
3. If \mathcal{A} and \mathcal{B} are sentences, then $(\mathcal{A} \wedge \mathcal{B})$ is a sentence.
4. If \mathcal{A} and \mathcal{B} are sentences, then $(\mathcal{A} \vee \mathcal{B})$ is a sentence.
5. If \mathcal{A} and \mathcal{B} are sentences, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a sentence.
6. If \mathcal{A} and \mathcal{B} are sentences, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a sentence.
7. Nothing else is a sentence.

Definitions like this are called *recursive*. Recursive definitions begin with some specifiable base elements, and then present ways to generate indefinitely many more elements by compounding together previously established ones. To give you a better idea of what a recursive definition is, we can give a recursive definition of the idea of *an ancestor of mine*. We specify a base clause.

- › My parents are ancestors of mine.

and then offer further clauses like:

- › If x is an ancestor of mine, then x 's parents are ancestors of mine.
- › Nothing else is an ancestor of mine.

Using this definition, we can easily check to see whether someone is my ancestor: just check whether she is the parent of the parent of ... one of my parents. And the same is true for our recursive definition of sentences of *Sentential*. Just as the recursive definition allows complex sentences to be built up from simpler parts, the definition allows us to decompose sentences into their simpler parts. And if we get down to atomic sentences, then we are ok.

Let's consider some examples.

Suppose we want to know whether or not ' $\neg\neg\neg D$ ' is a sentence of *Sentential*. Looking at the second clause of the definition, we know that ' $\neg\neg\neg D$ ' is a sentence if ' $\neg\neg D$ ' is a sentence. So now we need to ask whether or not ' $\neg\neg D$ ' is a sentence. Again looking at the second clause of the definition, ' $\neg\neg D$ ' is a sentence if ' $\neg D$ ' is. Again, ' $\neg D$ ' is a sentence if ' D ' is a sentence. Now ' D ' is an atomic sentence of *Sentential*, so we know that ' D ' is a sentence by the first clause of the definition. So for a compound sentence like ' $\neg\neg\neg D$ ', we must apply the definition repeatedly. Eventually we arrive at the atomic sentences from which the sentence is built up.

Next, consider the example ' $\neg(P \wedge \neg(\neg Q \vee R))$ '. Looking at the second clause of the definition, this is a sentence if ' $(P \wedge \neg(\neg Q \vee R))$ ' is. And this is a sentence if *both* ' P ' and ' $\neg(\neg Q \vee R)$ ' are sentences. The former is an atomic sentence, and the latter is a sentence if ' $\neg Q \vee R$ ' is a sentence. It is. Looking at the fourth clause of the definition, this is a sentence if both ' $\neg Q$ ' and ' R ' are sentences. And both are!

Ultimately, every sentence is constructed nicely out of atomic sentences. When we are dealing with a *sentence* other than an atomic sentence, we can see that there must be some sentential connective that was introduced *last*, when constructing the sentence. We call that the **MAIN CONNECTIVE** of the sentence. In the case of ' $\neg\neg\neg D$ ', the main connective is the very first ' \neg ' sign. In the case of ' $(P \wedge \neg(\neg Q \vee R))$ ', the main connective is ' \wedge '. In the case of ' $((\neg E \vee F) \rightarrow \neg\neg G)$ ', the main connective is ' \rightarrow '.

The recursive structure of sentences in *Sentential* 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. We will return to this point in chapter 3.

The recursive structure of sentences in *Sentential* also allows us to give a formal definition of the *scope* of a negation (mentioned in §5.2). The scope of a ' \neg ' is the subsentence for which ' \neg ' is the main connective. So in a sentence like:

$$(P \wedge (\neg(R \wedge B) \leftrightarrow Q))$$

this was constructed by conjoining ' P ' with ' $(\neg(R \wedge B) \leftrightarrow Q)$ '. This last sentence was constructed by placing a biconditional between ' $\neg(R \wedge B)$ ' and ' Q '. And the former of these sentences – a subsentence of our original sentence – is a sentence for which ' \neg ' is the main connective. So the scope of the negation is just ' $\neg(R \wedge B)$ '. More generally:

The SCOPE of an instance of a connective (in a sentence) is the subsentence which that has that instance of the connective as its main connective.

I talk of ‘instances’ of a connective because, in an example like ‘ $\neg\neg A$ ’, there are two occurrences of the negation connective, with different scopes – one is the main connective of the whole sentence, the other has just ‘ $\neg A$ ’ as its scope.

6.3 Bracketing conventions

Strictly speaking, the brackets in ‘ $(Q \wedge R)$ ’ are an indispensable part of the sentence. Part of this is because we might use ‘ $(Q \wedge R)$ ’ as a subsentence in a more complicated sentence. For example, we might want to negate ‘ $(Q \wedge R)$ ’, obtaining ‘ $\neg(Q \wedge R)$ ’. If we just had ‘ $Q \wedge R$ ’ without the brackets and put a negation in front of it, we would have ‘ $\neg Q \wedge R$ ’. It is most natural to read this as meaning the same thing as ‘ $(\neg Q \wedge R)$ ’. But as we saw in §5.2, this is very different from ‘ $\neg(Q \wedge R)$ ’.

Strictly speaking, then, ‘ $Q \wedge R$ ’ is *not* a sentence. It is a mere *expression*.

When working with Sentential, however, it will make our lives easier if we are sometimes a little less than strict. So, here are some convenient conventions.

First, we allow ourselves to omit the *outermost* brackets of a sentence. Thus we allow ourselves to write ‘ $Q \wedge R$ ’ instead of the sentence ‘ $(Q \wedge R)$ ’. However, we must remember to put the brackets back in, when we want to embed the sentence into a more complicated sentence!

Second, it can be a bit painful to stare at long sentences with many nested pairs of brackets. To make things a bit easier on the eyes, we shall allow ourselves to use brackets in varied sizes, which sometimes helps in seeing which brackets pair up with each other.

Combining these two conventions, we can rewrite the unwieldy sentence

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

rather more readably as follows:

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

The scope of each connective is now much clearer.

Practice exercises

A. For each of the following: (a) Is it a sentence of Sentential, strictly speaking? (b) Is it a sentence of Sentential, allowing for our relaxed bracketing conventions?

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

B. Are there any sentences of Sentential that contain no atomic sentences? Explain your answer.

C. What is the scope of each connective in the sentence

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

7 Use and mention

In this chapter, I have talked a lot *about* sentences. So I need to pause to explain an important, and very general, point.

7.1 Quotation conventions

Consider these two sentences:

- › Malcolm Turnbull is the Prime Minister.
- › The expression ‘Malcolm Turnbull’ is composed of two upper case letters and thirteen lower case letters.

When we want to talk about the Prime Minister, we *use* his name. When we want to talk about the Prime Minister’s name, we *mention* that name. And we do so by putting it in quotation marks.

There is a general point here. When we want to talk about things in the world, we just *use* words. When we want to talk about words, we typically have to *mention* those words.¹ We need to indicate that we are mentioning them, rather than using them. To do this, some convention is needed. We can put them in quotation marks, or display them centrally in the page (say). So this sentence:

- › ‘Malcolm Turnbull’ is the Prime Minister.

says that some *expression* is the Prime Minister. And that’s false. The *man* is the Prime Minister; his *name* isn’t. Conversely, this sentence:

- › Malcolm Turnbull is composed of two upper case letters and thirteen lower case letters.

¹ More generally, when we want to talk about something we use its name. So when we want to talk about an expression, we use the name of the expression – which is just the expression enclosed in quotation marks. Mentioning an expression is using a name of that expression.

also says something false: Malcolm Turnbull is a man, made of meat rather than letters. One final example:

- › ‘‘Malcolm Turnbull’’ is the name of ‘Malcolm Turnbull’.

On the left-hand-side, here, we have the name of a name. On the right hand side, we have a name. Perhaps this kind of sentence only occurs in logic textbooks, but it is true.

Those are just general rules for quotation, and you should observe them carefully in all your work! To be clear, the quotation-marks here do not indicate indirect speech. They indicate that you are moving from talking about an object, to talking about the name of that object.

7.2 Object language and metalanguage

These general quotation conventions are of particular importance for us. After all, we are describing a formal language here, *Sentential*, and so we are often *mentioning* expressions from *Sentential*.

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.

For the most part, the object language in this chapter has been the formal language that we have been developing: *Sentential*. The metalanguage is English. Not conversational English exactly, but English supplemented with some additional vocabulary which helps us to get along.

Now, I have used italic upper case letters for atomic sentences of *Sentential*:

$$A, B, C, Z, A_1, B_4, A_{25}, J_{375}, \dots$$

These are sentences of the object language (*Sentential*). They are not sentences of English. So I must not say, for example:

- › *D* is an atomic sentence of *Sentential*.

Obviously, I am trying to come out with an English sentence that says something about the object language (*Sentential*). But ‘*D*’ is a sentence of *Sentential*, and no part of English. So the preceding is gibberish, just like:

- › Schnee ist weiß is a German sentence.

What we surely meant to say, in this case, is:

- › ‘Schnee ist weiß’ is a German sentence.

Equally, what we meant to say above is just:

- › ‘*D*’ is an atomic sentence of *Sentential*.

The general point is that, whenever we want to talk in English about some specific expression of *Sentential*, we need to indicate that we are *mentioning* the expression, rather than *using* it. We can either deploy quotation marks, or we can adopt some similar convention, such as placing it centrally in the page.

7.3 Script fonts, and recursive definitions revisited

However, we do not just want to talk about *specific* expressions of *Sentential*. We also want to be able to talk about *any arbitrary* sentence of *Sentential*. Indeed, I had to do this in §6, when I presented the recursive definition of a sentence of *Sentential*. I used upper case script font letters to do this, namely:

A, B, C, D, ...

These symbols do not belong to *Sentential*. Rather, they are part of our (augmented) metalanguage that we use to talk about *any* expression of *Sentential*. To repeat the second clause of the recursive definition of a sentence of *Sentential*, we said:

3. If *A* is a sentence, then $\neg A$ is a sentence.

This talks about *arbitrary* sentences. If we had instead offered:

- › If ‘*A*’ is a sentence, then ‘ $\neg A$ ’ is a sentence.

this would not have allowed us to determine whether ‘ $\neg B$ ’ is a sentence. To emphasise, then:

‘*A*’ is a symbol in augmented English, which we use to talk about any *Sentential* expression. ‘*A*’ is a particular atomic sentence of *Sentential*.

To come at this distinction a slightly different way, while ‘*A*’ designates a sentence of *Sentential*, it can designate a different sentence on different occasions. It behaves a bit like a pronoun – ‘it’ always designates some object, but a different one in different circumstances of use. By contrast, ‘*A*’ always names just that atomic sentence of *Sentential* which is just the first letter of the English alphabet.

This last example raises a further complication for our quotation conventions. I have not included any quotation marks in the clauses of our recursive definition of a sentence of Sentential in §6.2. Should I have done so?

The problem is that the expression on the right-hand-side of most of our recursive clauses are not sentences of English, since they contain Sentential connectives, like ‘ \neg ’. Consider clause 2. We might try to write:

2'. If \mathcal{A} is a sentence, then ‘ $\neg\mathcal{A}$ ’ is a sentence.

But this is no good: ‘ $\neg\mathcal{A}$ ’ is not a Sentential sentence, since ‘ \mathcal{A} ’ is a symbol of (augmented) English rather than a symbol of Sentential. What we really want to say is something like this:

2''. If \mathcal{A} is a sentence, then the expression that consists of the symbol ‘ \neg ’, followed immediately by the sentence \mathcal{A} , is also a sentence.

This is impeccable, but rather long-winded. But we can avoid long-windedness by creating our own conventions. We can perfectly well stipulate that an expression like ‘ $\neg\mathcal{A}$ ’ should simply be read as abbreviating the long-winded account. So, *officially*, the metalanguage expression ‘ $\neg\mathcal{A}$ ’ simply abbreviates:

the expression that consists of the symbol ‘ \neg ’ followed by the sentence \mathcal{A}

and similarly, for expressions like ‘ $(\mathcal{A} \wedge \mathcal{B})$ ’, ‘ $(\mathcal{A} \vee \mathcal{B})$ ’, etc. The latter is the expression which consists of the sentence \mathcal{A} followed by the symbol ‘ \vee ’ followed by the sentence \mathcal{B} .

7.4 Quotation conventions for arguments

One of our main purposes for using Sentential is to study arguments, and that will be our concern in chapter 3. In English, the premises of an argument are often expressed by individual sentences, and the conclusion by a further sentence. Since we can symbolise English sentences, we can symbolise English arguments using Sentential. Thus we might ask whether the argument whose premises are the Sentential sentences ‘ \mathcal{A} ’ and ‘ $\mathcal{A} \rightarrow \mathcal{C}$ ’, and whose conclusion is the Sentential sentence ‘ \mathcal{C} ’, is valid. However, it is quite a mouthful to write that every time. So instead I shall introduce another bit of abbreviation. This:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$$

abbreviates:

the argument with premises $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ and conclusion \mathcal{C}

To avoid unnecessary clutter, we shall not regard this as requiring quotation marks around it. This is a name of an argument, not an argument itself. (Note, then, that ‘ \therefore ’ is a symbol of our augmented *metalanguage*, and not a new symbol of Sentential.)

Practice exercises

A. For each of the following: Are the quotation marks correctly used, strictly speaking? If not, propose a corrected version.

1. Snow is not a sentence of English.
2. ' $\mathcal{A} \rightarrow \mathcal{C}$ ' is a sentence of Sentential.
3. ' $\neg\mathcal{A}$ ' is the expression consisting of the symbol ' \neg ' followed by the upper case script letter ' \mathcal{A} '.
4. If ' \mathcal{A} ' is a sentence, so is ' $(\mathcal{A} \vee \mathcal{A})$ '.
5. ' \mathcal{A} ' has the same number of characters as ' A '.

Chapter 3

Truth tables

8 Schematic truth tables

So much for the grammar of *Sentential*. We turn now to the intended interpretation of *Sentential* sentences.

We have seen how to use a symbolisation key in §4.1 to temporarily assign a meaning to some of the atomic sentences of *Sentential*. This symbolisation key will at the very least assign a truth-value to those atomic sentences. Any non-atomic sentence of *Sentential* is composed of atomic sentences with sentential connectives. It turns out that we don't need to know anything more about the atomic sentences of *Sentential* than their truth values to assign a truth value to those non-atomic, or COMPOUND, sentences. More generally, the truth value of any compound sentence depends only on the truth value of the subsentences that comprise it. In order to know the truth value of ' $(D \wedge E)$ ', for instance, you only need to know the truth value of ' D ' and the truth value of ' E '. In order to know the truth value of ' $((D \wedge E) \vee F)$ ', you need only know the truth value of ' $(D \wedge E)$ ' and ' F '. And so on. This is fact a good part of the reason why we've chose these connectives, and chose their English near-equivalents as our structural words.

We introduced five connectives in chapter 2. So we simply need to explain how they map between truth values. For convenience, we shall represent 'True' by 'T' and 'False' by 'F'. (But just to be clear, the two truth values are True and False; the truth values are not *letters*!)

Not only does the truth value of a compound sentence of *Sentential* depend only on the truth values assigned to its subsentences, it does so *uniquely*. We may happily refer to *the* truth value of ' \mathcal{A} ' as determined by its constituents. This means we may represent the pattern of dependence of compound sentence truth values on subsentence truth values in a simple table.

Negation. 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. We can summarize this dependence in the *schematic truth table* for negation:

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

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

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

Note that conjunction is *symmetrical*. The truth value for ' $\mathcal{A} \wedge \mathcal{B}$ ' is always the same as the truth value for ' $\mathcal{B} \wedge \mathcal{A}$ '.

Disjunction. Recall that ' \vee ' always represents inclusive or. So, for any sentences \mathcal{A} and \mathcal{B} , ' $\mathcal{A} \vee \mathcal{B}$ ' is true if and only if either \mathcal{A} or \mathcal{B} is true. We can summarize this in the schematic 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.

Conditional. I'm just going to come clean and admit it. Conditionals are a problem in Sentential. This is not because there is any problem finding a truth table for the connective ' \rightarrow ', but rather because the truth table we put forward seems to make ' \rightarrow ' behave in ways that are different to the way that the English counterpart 'if ... then ...' behaves. Exactly how much of a problem this poses is a matter of *philosophical* contention. I shall discuss a few of the subtleties in §§9.3 and 11.5.

We know at least this much: if \mathcal{A} is true and \mathcal{B} is false, then ' $\mathcal{A} \rightarrow \mathcal{B}$ ' should be false. For now, I am going to stipulate that this is the only type of case in which ' $\mathcal{A} \rightarrow \mathcal{B}$ ' is false. We can summarize this with a schematic truth table for the Sentential 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 $\mathcal{A} \rightarrow \mathcal{B}$ has a very different truth table from $\mathcal{B} \rightarrow \mathcal{A}$.

Biconditional. Since a biconditional is to be the same as the conjunction of a conditional running in each direction, we shall want the truth table for the biconditional to be:

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

Unsurprisingly, the biconditional is symmetrical.

9 Truth-functional connectives

9.1 The idea of truth-functionality

I want to introduce an important idea about sentence connectives.

A connective is TRUTH-FUNCTIONAL iff the truth value of a sentence with that connective as its main connective is uniquely determined by the truth value(s) of the constituent sentence(s).

We just saw that every connective in *Sentential* is truth-functional - we would not have been able to draw up our schematic truth tables if they had not been. The truth value of a negation is uniquely determined by the truth value of the unnegated sentence. The truth value of a conjunction is uniquely determined by the truth value of both conjuncts. The truth value of a disjunction is uniquely determined by the truth value of both disjuncts. And so on. To determine the truth value of some *Sentential* sentence, we only need to know the truth value of its components. This is why the study of *Sentential* is termed *truth-functional logic*.

In plenty of languages there are connectives that are not truth-functional. In English, for example, we can form a new sentence from any simpler sentence by prefixing it with 'It is necessarily the case that ...'. The truth value of this new sentence is not fixed solely by the truth value of the original sentence. For consider two true sentences:

1. $2 + 2 = 4$.
2. Shostakovich wrote fifteen string quartets.

Whereas it is necessarily the case that $2 + 2 = 4$,¹ it is not *necessarily* the case that Shostakovich wrote fifteen string quartets. If Shostakovich had died earlier, he would have failed to finish Quartet no. 15; if he had lived longer, he might have written a few more. So 'It is necessarily the case that ...' is a connective of English, but it is not *truth-functional*.

¹ Given that the English numeral '2' names the number two, and the numeral '4' names the number four, and '+' names addition, then it must be that the result of adding two to itself is four. This is not to say that '2' had to be used in the way we actually use it - if '2' had named three, the sentence would have been false. But in its actual use, it is a necessary truth.

9.2 Symbolising versus translating

All of the connectives of *Sentential* are truth-functional. But more than that: they really do nothing *but* map us between truth values.

When we symbolise a sentence or an argument in *Sentential*, we ignore everything *besides* the contribution that the truth values of a component might make to the truth value of the whole. There are subtleties to our ordinary claims that far outstrip their mere truth values. Sarcasm; poetry; snide implicature; emphasis; these are important parts of everyday discourse. But none of this is retained in *Sentential*. As remarked in §5, *Sentential* cannot capture the subtle differences between the following English sentences:

1. Jon is fat and Jon is quick
2. Although Jon is fat, Jon is quick
3. Despite being fat, Jon is quick
4. Jon is quick, albeit fat
5. Jon's fatness notwithstanding, he is quick

All of the above sentences will be symbolised with the same *Sentential* sentence, perhaps ' $F \wedge Q$ '.

I keep saying that we use *Sentential* sentences to *symbolise* English sentences. Many other textbooks talk about *translating* English sentences into *Sentential*. But a good translation should preserve certain facets of meaning, and – as I have just pointed out – *Sentential* just cannot do that. This is why I shall speak of *symbolising* English sentences, rather than of *translating* them.

This affects how we should understand our symbolisation keys. Consider a key like:

- F : Jon is fat.
 Q : Jon is quick.

Other textbooks will understand this as a stipulation that the *Sentential* sentence ' F ' should *mean* that Jon is fat, and that the *Sentential* sentence ' Q ' should *mean* that Jon is quick. But *Sentential* just is totally unequipped to deal with *meaning*. The preceding symbolisation key is doing no more nor less than stipulating that the *Sentential* sentence ' F ' should take the same truth value as the English sentence 'Jon is fat' (whatever that might be), and that the *Sentential* sentence ' Q ' should take the same truth value as the English sentence 'Jon is quick' (whatever that might be).

When we treat a *Sentential* sentence as *symbolising* an English sentence, we are stipulating that the *Sentential* sentence is to take the same truth value as that English sentence.

9.3 Indicative versus subjunctive conditionals

I want to bring home the point that Sentential can *only* deal with truth functions by considering the case of the conditional. When I introduced the schematic truth table for the material conditional in §8, I did not say much to justify it. Let me now offer a justification, which follows Dorothy Edgington.²

Suppose that Lara has drawn some shapes on a piece of paper, and coloured some of them in. I have not seen them, but I claim:

If any shape is grey, then that shape is also circular.

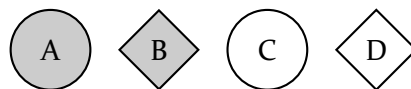
As it happens, Lara has drawn the following:



In this case, my claim is surely true. Shapes C and D are not grey, and so can hardly present *counterexamples* to my claim. Shape A is grey, but fortunately it is also circular. So my claim has no counterexamples. It must be true. And that means that each of the following *instances* of my claim must be true too:

- › If A is grey, then it is circular (true antecedent, true consequent)
- › If C is grey, then it is circular (false antecedent, true consequent)
- › If D is grey, then it is circular (false antecedent, false consequent)

However, if Lara had drawn a fourth shape, thus:



then my claim would have been false. So it must be that this claim is false:

- › If B is grey, then it is a circular (true antecedent, false consequent)

² Dorothy Edgington, 'Conditionals', 2006, in the *Stanford Encyclopedia of Philosophy* (<http://plato.stanford.edu/entries/conditionals/>).

Now, recall that every connective of *Sentential* has to be truth-functional. This means that the mere truth value of the antecedent and consequent must uniquely determine the truth value of the conditional as a whole. Thus, from the truth values of our four claims – which provide us with all possible combinations of truth and falsity in antecedent and consequent – we can read off the truth table for the material conditional.

What this argument shows is that ‘ \rightarrow ’ is the *only* candidate for a truth-functional conditional. Otherwise put, *it is the best conditional that Sentential can provide*. But is it any good, as a surrogate for the conditionals we use in everyday language? Consider two sentences:

1. If Mitt Romney had won the 2012 election, then he would have been the 45th President of the USA.
2. If Mitt Romney had won the 2012 election, then he would have turned into a helium-filled balloon and floated away into the night sky.

Sentence **1** is true; sentence **2** is false. But both have false antecedents and false consequents. So the truth value of the whole sentence is not uniquely determined by the truth value of the parts. Do not just blithely assume that you can adequately symbolise an English ‘if ..., then ...’ with *Sentential*’s ‘ \rightarrow ’.

The crucial point is that sentences **1** and **2** employ *subjunctive* conditionals, rather than *indicative* conditionals. They ask us to imagine something contrary to fact – Mitt Romney lost the 2012 election – and then ask us to evaluate what *would* have happened in that case. The classic illustration of the difference between the indicative and subjunctive conditional comes from pairs like these:

3. If a dingo *didn’t take* Azaria Chamberlain, something else *did*.
4. If a dingo *hadn’t taken* Azaria Chamberlain, something else *would have*.

The indicative conditional in **3** is true, given the actual historical fact that she was taken. But is the subjunctive, or *counterfactual*, in **4** also true? It seems not. (She was not destined to be taken by something or other, and if the dingo hadn’t intervened, she wouldn’t have disappeared at all.)

The point to take away from this is that subjunctive conditionals cannot be tackled using ‘ \rightarrow ’. This is not to say that they cannot be tackled by any formal logical language, only that *Sentential* is not up to the job. Only indicative conditionals are able to be modelled by *Sentential*.

In fact there are even difficulties with indicatives in *Sentential*. I shall say a little more about those difficulties in §11.5 and in §27.5. For now, I shall content myself with the observation that ‘ \rightarrow ’ is the only plausible candidate for a truth-functional conditional, but that many English conditionals cannot be represented adequately using ‘ \rightarrow ’. *Sentential* is an intrinsically limited language. But this is only a problem if you try to use it to do things it wasn’t designed to do.

10 Complete truth tables

So far, we have considered assigning truth values to *Sentential* sentences indirectly. We have said, for example, that a *Sentential* sentence such as ‘*B*’ is to take the same truth value as the English sentence ‘Big Ben is in London’ (whatever that truth value may be). But we can also assign truth values *directly*. We can simply stipulate that ‘*B*’ is to be true, or stipulate that it is to be false.

A VALUATION is any assignment of truth values to some atomic sentences of *Sentential*. It assigns exactly one truth value, either True or False, to each of the sentences in question.

A valuation is a temporary assignment of ‘meanings’ to *Sentential* sentences, in much the same way as a symbolisation key might be. What is distinctive about *Sentential* is that almost all of its basic vocabulary – the atomic sentences – only get their meanings in this temporary fashion. The only parts of *Sentential* that get their meanings permanently are the connectives, which always have a fixed interpretation. This is rather unlike English, where most words have their meanings on a permanent basis. But there are some words in English – like pronouns (‘he’, ‘she’) and demonstratives (‘this’, ‘that’) – that get their meaning assigned temporarily, and then can be reused with a different meaning in another context. Of course we don’t have anything so explicit and deliberate as a valuation or a symbolisation key in English to assign a meaning to a particular use of ‘this’.

We introduced schematic truth tables in §8. These showed what truth value a compound sentence with a certain structure was determined to have by the truth values of its subsentences, whatever they might be. We now introduce a closely related idea, that of a truth table. This shows how a specific compound sentence has its truth value determined by the truth values of its specific *atomic subsentences*, across all the possible ways that those atomic subsentences might be assigned True and False.

You will no doubt have realised that a way of assigning True and False to atomic sentences is a valuation. So we can say: a TRUTH TABLE summarises how the truth value of a compound sentence depends on the possible valuations of its atomic subsentences. Each row of a truth table represents a possible valuation. The entire truth table represents all possible valuations. And the truth table provides us with a means to calculate the truth value of complex sentences, on each possible valuation. This is pretty abstract. So it might be easiest to explain with an example.

10.1 A worked example

Consider the sentence $(H \wedge I) \rightarrow H$. There are four possible ways to assign True and False to the atomic sentence $'H'$ and $'I'$ – four possible valuations – which we can represent as follows:

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

To calculate the truth value of the entire sentence $(H \wedge I) \rightarrow H$, we first copy the truth values for the atomic sentences and write them underneath the letters in the sentence:

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

Now consider the subsentence $(H \wedge I)$. This is a conjunction, $(\mathcal{A} \wedge \mathcal{B})$, with $'H'$ as \mathcal{A} and with $'I'$ as \mathcal{B} . The schematic truth table for conjunction gives the truth conditions for *any* sentence of the form $(\mathcal{A} \wedge \mathcal{B})$, whatever \mathcal{A} and \mathcal{B} might be. It summarises the point that a conjunction is true iff both conjuncts are true. In this case, our conjuncts are just $'H'$ and $'I'$. They are both true on (and only on) the first line of the truth table. Accordingly, we can calculate the truth value of the conjunction on all four rows.

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

Now, the entire sentence that we are dealing with is a conditional, $\mathcal{A} \rightarrow \mathcal{B}$, with $(H \wedge I)$ as \mathcal{A} and with $'H'$ as \mathcal{B} . On the second row, for example, $(H \wedge I)$ 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:

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

The conditional is the main logical connective of the sentence. And the column of ‘T’s underneath the conditional tells us that the sentence ‘ $(H \wedge I) \rightarrow H$ ’ 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. Since we have considered all four possible assignments of truth and falsity to ‘ H ’ and ‘ I ’ – since, that is, we have considered all the different *valuations* – we can say that ‘ $(H \wedge I) \rightarrow H$ ’ is true on every valuation.

In this example, I have not repeated all of the entries in every column 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	$(H \wedge I) \rightarrow H$
T	T	T T T T
T	F	T F F T
F	T	F F T T
F	F	F F T F

Most of the columns underneath the sentence are only there for bookkeeping purposes. The column that matters most is the column underneath the *main connective* for the sentence, since this tells you the truth value of the entire sentence. I have emphasised this, by putting this column in bold. When you work through truth tables yourself, you should similarly emphasise it (perhaps by underlining).

10.2 Building complete truth tables

A COMPLETE TRUTH TABLE has a line for every possible assignment of True and False to the relevant atomic sentences. Each line represents a *valuation*, and a complete truth table has a line for all the different valuations.

The size of the complete truth table depends on the number of different atomic sentences in the table. A sentence that contains only one atomic sentence requires only two rows, as in the schematic truth table for negation. This is true even if the same letter is repeated many times, as in the sentence ‘ $((C \leftrightarrow C) \rightarrow C) \wedge \neg(C \rightarrow C)$ ’. The complete truth table requires only two lines because there are only two possibilities: ‘ C ’ can be true or it can be false. The truth table for this sentence looks like this:

C	$((C \leftrightarrow C) \rightarrow C) \wedge \neg(C \rightarrow C)$
T	T T T T T F F T T T
F	F T 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., the sentence is false regardless of whether ' C ' is true or false. It is false on every valuation.

A sentence that contains two atomic sentences requires four lines for a complete truth table, as in the schematic truth tables, and as in the complete truth table for ' $(H \wedge I) \rightarrow H$ '.

A sentence that contains three atomic sentences requires eight lines:

M	N	P	$M \wedge (N \vee P)$
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 ' $M \wedge (N \vee P)$ ' can 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 atomic sentences 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 atomic sentences, then it must have 2^n lines.¹

In order to fill in the columns of a complete truth table, begin with the right-most atomic sentence and alternate between 'T' and 'F'. In the next column to the left, write two 'T's, write two 'F's, and repeat. For the third atomic sentence, write four 'T's followed by four 'F's. This yields an eight line truth table like the one above. For a 16 line truth table, the next column of atomic sentences should have eight 'T's followed by eight 'F's. For a 32 line table, the next column would have 16 'T's followed by 16 'F's. And so on.

¹ Since the values of atomic sentences are independent of each other, each new atomic sentence \mathcal{A}_{n+1} we consider is capable of being true or false on every existing valuation on $\mathcal{A}_1, \dots, \mathcal{A}_n$, and so there must be twice as many valuations on $\mathcal{A}_1, \dots, \mathcal{A}_n, \mathcal{A}_{n+1}$ as on $\mathcal{A}_1, \dots, \mathcal{A}_n$.

10.3 More bracketing conventions

Consider these two sentences:

$$\begin{aligned} &((A \wedge B) \wedge C) \\ &(A \wedge (B \wedge C)) \end{aligned}$$

These have the same truth table. Consequently, it will never make any difference from the perspective of truth value – which is all that Sentential cares about (see §9) – which of the two sentences we assert (or deny). And since the order of the brackets does not matter, I shall allow us to drop them. In short, we can save some ink and some eyestrain by writing:

$$A \wedge B \wedge C$$

The general point is that, if we just have a long list of conjunctions, we can drop the inner brackets. (I already allowed us to drop outermost brackets in §6.) The same observation holds for disjunctions. Since the following sentences have exactly the same truth table:

$$\begin{aligned} &((A \vee B) \vee C) \\ &(A \vee (B \vee C)) \end{aligned}$$

we can simply write:

$$A \vee B \vee C$$

And generally, if we just have a long list of disjunctions, we can drop the inner brackets. *But be careful.* These two sentences have *different* truth tables:

$$\begin{aligned} &((A \rightarrow B) \rightarrow C) \\ &(A \rightarrow (B \rightarrow C)) \end{aligned}$$

So if we were to write:

$$A \rightarrow B \rightarrow C$$

it would be dangerously ambiguous. So we must not do the same with conditionals. Equally, these sentences have different truth tables:

$$\begin{aligned} &((A \vee B) \wedge C) \\ &(A \vee (B \wedge C)) \end{aligned}$$

So if we were to write:

$$A \vee B \wedge C$$

it would be dangerously ambiguous. *Never write this.* The moral is: you can drop brackets when dealing with a long list of conjunctions, or when dealing with a long list of disjunctions. But that's it.

Practice exercises

A. Offer complete truth tables for each of the following:

1. $A \rightarrow A$
2. $C \rightarrow \neg C$
3. $(A \leftrightarrow B) \leftrightarrow \neg(A \leftrightarrow \neg B)$
4. $(A \rightarrow B) \vee (B \rightarrow A)$
5. $(A \wedge B) \rightarrow (B \vee A)$
6. $\neg(A \vee B) \leftrightarrow (\neg A \wedge \neg B)$
7. $((A \wedge B) \wedge \neg(A \wedge B)) \wedge C$
8. $((A \wedge B) \wedge C) \rightarrow B$
9. $\neg((C \vee A) \vee B)$

B. Check all the claims made in introducing the new notational conventions in §10.3, i.e., show that:

1. ' $((A \wedge B) \wedge C)$ ' and ' $(A \wedge (B \wedge C))$ ' have the same truth table
2. ' $((A \vee B) \vee C)$ ' and ' $(A \vee (B \vee C))$ ' have the same truth table
3. ' $((A \vee B) \wedge C)$ ' and ' $(A \vee (B \wedge C))$ ' do not have the same truth table
4. ' $(A \rightarrow B) \rightarrow C$ ' and ' $(A \rightarrow (B \rightarrow C))$ ' do not have the same truth table

Also, check whether:

5. ' $(A \leftrightarrow B) \leftrightarrow C$ ' and ' $(A \leftrightarrow (B \leftrightarrow C))$ ' have the same truth table

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

11 Semantic concepts

In the previous section, we introduced the idea of a valuation and showed how to determine the truth value of any Sentential sentence, on any valuation, using a truth table. In this section, we shall introduce some related ideas, and show how to use truth tables to test whether or not they apply.

11.1 Tautologies and contradictions

In §3, I explained *necessary truth* and *necessary falsity*. Both notions have close but imperfect surrogates in Sentential. We shall start with a surrogate for necessary truth.

\mathcal{A} is a TAUTOLOGY iff it is true on every valuation.

We can determine whether a sentence is a tautology just by using truth tables. If the sentence is true on every line of a complete truth table, then it is true on every valuation, so it is a tautology. In the example of §10, $(H \wedge I) \rightarrow H$ is a tautology.

This is only, though, a surrogate for necessary truth. There are some necessary truths that we cannot adequately symbolise in Sentential. An example is $2 + 2 = 4$. This *must* be true, but if we try to symbolise it in Sentential, the best we can offer is an atomic sentence, and no atomic sentence is a tautology.¹ Still, if we can adequately symbolise some English sentence using a Sentential sentence which is a tautology, then that English sentence expresses a necessary truth.

We have a similar surrogate for necessary falsity:

\mathcal{A} is a CONTRADICTION iff it is false on every valuation.

We can determine whether a sentence is a contradiction just by using truth tables. If the sentence is false on every line of a complete truth table, then it is false on every valuation, so it is a contradiction. In the example of §10, $((C \leftrightarrow C) \rightarrow C) \wedge \neg(C \rightarrow C)$ is a contradiction.

¹ At this risk of repeating myself: $2 + 2 = 4$ is necessarily true, but it is not necessarily true *in virtue of its structure*. A necessary truth is true, *with its actual meaning*, in every possible situation. A Sentential-tautology is true *in the actual situation* on every possible way of interpreting its atomic sentences. These are interestingly different notions.

11.2 Equivalence

Here is a similar, useful notion:

\mathcal{A} and \mathcal{B} are EQUIVALENT iff they have the same truth value on every valuation.

We have already made use of this notion, in effect, in §10.3; the point was that $(A \wedge B) \wedge C$ and $A \wedge (B \wedge C)$ are equivalent. Again, it is easy to test for equivalence using truth tables. Consider the sentences $\neg(P \vee Q)$ and $\neg P \wedge \neg Q$. Are they equivalent? To find out, we construct a truth table.

P	Q	$\neg(P \vee Q)$	$\neg P \wedge \neg Q$
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 false. On the final row, both are true. Since they match on every row, the two sentences are equivalent.

11.3 Consistency

In §3, I said that sentences are jointly consistent iff it is possible for all of them to be true at once. We can offer a surrogate for this notion too:

$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are JOINTLY CONSISTENT iff there is some valuation which makes them all true.

Derivatively, sentences are jointly inconsistent if there is no valuation that makes them all true. Again, it is easy to test for joint consistency using truth tables.

11.4 Entailment and validity

The following idea is closely related to that of joint consistency:

The sentences $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ ENTAIL the sentence \mathcal{C} if there is no valuation of the atomic sentences which makes all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ true and \mathcal{C} false.

Again, it is easy to test this with a truth table. Let us check whether $\neg L \rightarrow (J \vee L)$ and $\neg L$ entail J , we simply need to check whether there is any valuation which makes both $\neg L \rightarrow (J \vee L)$ and $\neg L$ true whilst making J false. So we use a truth table:

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

The only row on which both ' $\neg L \rightarrow (J \vee L)$ ' and ' $\neg L$ ' are true is the second row, and that is a row on which ' J ' is also true. So ' $\neg L \rightarrow (J \vee L)$ ' and ' $\neg L$ ' entail ' J '.

We now make an important observation:

If $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ entail \mathcal{C} , then $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$ is valid.

Here's why. If $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ entail \mathcal{C} , then there is no valuation which makes all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ true whilst making \mathcal{C} false. It is thus not possible – given the actual meanings of the connectives of **Sentential** – for the sentences $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ to jointly be true without \mathcal{C} being true too, so the argument $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$ is conclusive. Furthermore, because the conclusiveness of this argument doesn't depend on anything other than the structure of the sentences in the argument, it is also valid. (Because we consider every valuation, the conclusiveness of the argument cannot turn on the particular truth values assigned to the atomic sentences.)

In short, we have a way to test for the validity of English arguments. First, we symbolise them in **Sentential**, as having premises $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$, and conclusion \mathcal{C} . Then we test for entailment using truth tables.

11.5 The limits of these tests

We have reached an important milestone: a test for the validity of arguments! But, we should not get carried away just yet. It is important to understand the *limits* of our achievement. I shall illustrate these limits with three examples.

First, consider the argument:

1. Daisy is a small cow. So, Daisy is a cow.

To symbolise this argument in **Sentential**, we would have to use two different atomic sentences – perhaps ' S ' and ' C ' – for the premise and the conclusion respectively. Now, it is obvious that ' S ' does not entail ' C '. But the English argument surely seems valid – the structure of 'Daisy is a small cow' guarantees that 'Daisy is a cow' is true. (Note that a small cow might still be rather large, so we cannot fudge things by symbolising 'Daisy is a small cow' as a conjunction of 'Daisy is small' and 'Daisy is a cow'. We'll return to this sort of case in §16.) So our **Sentential**-based test for validity in English will have some *false*

negatives: it will classify some valid English arguments as invalid. But any argument that it classifies as valid will indeed be valid – it will not exhibit *false positives*.

Second, consider the sentence:

2. Jan is neither bald nor not-bald.

To symbolise this sentence in *Sentential*, we would offer something like ‘ $\neg J \wedge \neg\neg J$ ’. This a contradiction (check this with a truth-table). But sentence **2** does not itself seem like a contradiction; for we might have happily go on to add ‘Jan is on the borderline of baldness’!

Third, consider the following sentence:

3. It’s not the case that, if God exists, She answers malevolent prayers.

Symbolising this in *Sentential*, we would offer something like ‘ $\neg(G \rightarrow M)$ ’. Now, ‘ $\neg(G \rightarrow M)$ ’ entails ‘ G ’ (again, check this with a truth table). So if we symbolise sentence **3** in *Sentential*, it seems to entail that God exists. But that’s strange: surely even the atheist can accept sentence **3**, without contradicting herself! Some say that **3** would be better symbolised by ‘ $G \rightarrow \neg M$ ’, even though that doesn’t reflect the apparent form of the English sentence. ‘ $G \rightarrow \neg M$ ’ does not entail G . This symbolisation does a better job of reflecting the intuitive consequences of the English sentence **3**, but at the cost of abandoning a straightforward correspondence between the structure of English sentences and their *Sentential* symbolisations.

In different ways, these three examples highlight some of the limits of working with a language (like *Sentential*) that can *only* handle truth-functional connectives. Moreover, these limits give rise to some interesting questions in philosophical logic. The case of Jan’s baldness (or otherwise) raises the general question of what logic we should use when dealing with *vague* discourse. The case of the atheist raises the question of how to deal with the (so-called) *paradoxes of material implication*. Part of the purpose of this course is to equip you with the tools to explore these questions of *philosophical logic*. But we have to walk before we can run; we have to become proficient in using *Sentential*, before we can adequately discuss its limits, and consider alternatives. I will return one final time to the relation between ‘if’ and \rightarrow in §27.5.

11.6 The double-turnstile

We are going to use the notion of entailment rather a lot in this book. It will help us, then, to introduce a symbol that abbreviates it. Rather than saying that the *Sentential* sentences $\mathcal{A}_1, \mathcal{A}_2, \dots$ and \mathcal{A}_n together entail \mathcal{C} , we shall abbreviate this by:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vDash \mathcal{C}.$$

The symbol ‘ \vDash ’ is known as *the double-turnstile*, since it looks like a turnstile with two horizontal beams.

But let me be clear. ‘ \models ’ is not a symbol of **Sentential**. Rather, it is a symbol of our metalanguage, augmented English (recall the difference between object language and metalanguage from §7). So the metalanguage sentence:

$$4. P, P \rightarrow Q \models Q$$

is just an abbreviation for the English sentence:

$$5. \text{The Sentential sentences 'P' and 'P} \rightarrow \text{Q' entail 'Q'}$$

Note that there is no limit on the number of **Sentential** sentences that can be mentioned before the symbol ‘ \models ’. Indeed, we can even consider the limiting case:

$$6. \models C.$$

6 is false if there is a valuation which makes all the sentences appearing on the left hand side of ‘ \models ’ true and makes C false. Since *no* sentences appear on the left side of ‘ \models ’ in **6**, it is trivial to make ‘them’ all true. So it is false if there is a valuation which makes C false – and so **6** is true iff every valuation makes C true. Otherwise put, **6** says that C is a tautology. Equally:

$$7. \mathcal{A}_1, \dots, \mathcal{A}_n \models$$

says that $\mathcal{A}_1, \dots, \mathcal{A}_n$ are jointly inconsistent. It follows that

$$8. \mathcal{A} \models$$

says that \mathcal{A} is individually inconsistent. That is to say, \mathcal{A} is a contradiction.²

Here is an important connection between inconsistency and entailment.

$\mathcal{A}_1, \dots, \mathcal{A}_n \models C$ iff $\mathcal{A}_1, \dots, \mathcal{A}_n, \neg C \models$.

If every valuation which makes each of $\mathcal{A}_1, \dots, \mathcal{A}_n$ true also makes C true, then all of those valuations also make ‘ $\neg C$ ’ false. So there can be no valuation which makes each of $\mathcal{A}_1, \dots, \mathcal{A}_n, \neg C$ true. So those sentences are jointly inconsistent.

² If you find it difficult to see why ‘ $\models \mathcal{A}$ ’ should say that \mathcal{A} is a tautology, you should just take **6** as an abbreviation for that claim. Likewise you should take ‘ $\mathcal{A} \models$ ’ as abbreviating the claim that \mathcal{A} is a contradiction.

11.7 ‘ \vDash ’ versus ‘ \rightarrow ’

I now want to compare and contrast ‘ \vDash ’ and ‘ \rightarrow ’.

Observe: $\mathcal{A} \vDash \mathcal{C}$ iff there is no valuation of the atomic sentences that makes \mathcal{A} true and \mathcal{C} false.

Observe: $\mathcal{A} \rightarrow \mathcal{C}$ is a tautology iff there is no valuation of the atomic sentences that makes $\mathcal{A} \rightarrow \mathcal{C}$ false. Since a conditional is true except when its antecedent is true and its consequent false, $\mathcal{A} \rightarrow \mathcal{C}$ is a tautology iff there is no valuation that makes \mathcal{A} true and \mathcal{C} false.

Combining these two observations, we see that $\mathcal{A} \rightarrow \mathcal{C}$ is a tautology iff $\mathcal{A} \vDash \mathcal{C}$.³ But there is a really important difference between ‘ \vDash ’ and ‘ \rightarrow ’:

‘ \rightarrow ’ is a sentential connective of Sentential.
 ‘ \vDash ’ is a symbol of augmented English.

Indeed, when ‘ \rightarrow ’ is flanked with two Sentential sentences, the result is a longer Sentential sentence. By contrast, when we use ‘ \vDash ’, we form a metalinguistic sentence that *mentions* the surrounding Sentential sentences.

If ‘ $\mathcal{A} \rightarrow \mathcal{C}$ ’ is a tautology, then $\mathcal{A} \vDash \mathcal{C}$. But ‘ $\mathcal{A} \rightarrow \mathcal{C}$ ’ can be true on a valuation without being a tautology, and so can be true on a valuation even when \mathcal{A} doesn’t entail \mathcal{C} . Sometimes people are inclined to confuse entailment and conditionals, perhaps because they are tempted by the thought that we can only establish the truth of a conditional by *logically deriving* the consequent from the antecedent. But while this is the way to establish the truth of a tautologous conditional, most conditionals posit a weaker relation between antecedent and consequent than that – for example, a causal or statistical relationship might be enough to justify the truth of the conditional ‘If you smoke, then you’ll lower your life expectancy’.

Practice exercises

A. Revisit your answers to §10A. Determine which sentences were tautologies, which were contradictions, and which were neither tautologies nor contradictions.

B. Use truth tables to determine whether these sentences are jointly consistent, or jointly inconsistent:

³ This result sometimes goes under a fancy title: *the Deduction Theorem for Sentential*. It is easy to see that this more general result follows from the Deduction Theorem: $\mathcal{B}_1, \dots, \mathcal{B}_n, \mathcal{A} \vDash \mathcal{C}$ iff $\mathcal{B}_1, \dots, \mathcal{B}_n \vDash \mathcal{A} \rightarrow \mathcal{C}$.

1. $A \rightarrow A, \neg A \rightarrow \neg A, A \wedge A, A \vee A$
2. $A \vee B, A \rightarrow C, B \rightarrow C$
3. $B \wedge (C \vee A), A \rightarrow B, \neg(B \vee C)$
4. $A \leftrightarrow (B \vee C), C \rightarrow \neg A, A \rightarrow \neg B$

C. Use truth tables to determine whether each argument is valid or invalid.

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

D. Answer each of the questions below and justify your answer.

1. Suppose that \mathcal{A} and \mathcal{B} are equivalent. What can you say about $\mathcal{A} \leftrightarrow \mathcal{B}$?
2. Suppose that $(\mathcal{A} \wedge \mathcal{B}) \rightarrow \mathcal{C}$ is neither a tautology nor a contradiction. What can you say about whether $\mathcal{A}, \mathcal{B} \therefore \mathcal{C}$ is valid?
3. Suppose that \mathcal{A}, \mathcal{B} and \mathcal{C} are jointly inconsistent. What can you say about $(\mathcal{A} \wedge \mathcal{B} \wedge \mathcal{C})$?
4. Suppose that \mathcal{A} is a contradiction. What can you say about whether $\mathcal{A}, \mathcal{B} \vDash \mathcal{C}$?
5. Suppose that \mathcal{C} is a tautology. What can you say about whether $\mathcal{A}, \mathcal{B} \vDash \mathcal{C}$?
6. Suppose that \mathcal{A} and \mathcal{B} are equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?
7. Suppose that \mathcal{A} and \mathcal{B} are *not* equivalent. What can you say about $(\mathcal{A} \vee \mathcal{B})$?

E. Consider the following principle:

- › Suppose \mathcal{A} and \mathcal{B} are equivalent. Suppose an argument contains \mathcal{A} (either as a premise, or as the conclusion). The validity of the argument would be unaffected, if we replaced \mathcal{A} with \mathcal{B} .

Is this principle correct? Explain your answer.

12 Truth table shortcuts

With practice, you will quickly become adept at filling out truth tables. In this section, I want to give you some permissible shortcuts to help you along the way.

12.1 Working through truth tables

You will quickly find that you do not need to copy the truth value of each atomic sentence, but can simply refer back to them. So you can speed things up by writing:

P	Q	$(P \vee Q) \leftrightarrow \neg P$	
T	T	T	FF
T	F	T	FF
F	T	T	TT
F	F	F	FT

You also know for sure that a disjunction is true whenever one of the disjuncts is true. So if you find a true disjunct, there is no need to work out the truth values of the other disjuncts. Thus you might offer:

P	Q	$(\neg P \vee \neg Q) \vee \neg P$		
T	T	F	FF	FF
T	F	F	TT	TF
F	T			TT
F	F			TT

Equally, you know for sure that a conjunction is false whenever one of the conjuncts is false. So if you find a false conjunct, there is no need to work out the truth value of the other conjunct. Thus you might offer:

P	Q	$\neg(P \wedge \neg Q) \wedge \neg P$		
T	T			FF
T	F			FF
F	T	T	F	TT
F	F	T	F	TT

A similar short cut is available for conditionals. You immediately know that a conditional is true if either its consequent is true, or its antecedent is false. Thus you might present:

P	Q	$((P \rightarrow Q) \rightarrow P) \rightarrow P$		
T	T			T
T	F			T
F	T	T	F	T
F	F	T	F	T

So ' $((P \rightarrow Q) \rightarrow P) \rightarrow P$ ' is a tautology. In fact, it is an instance of *Peirce's Law*, named after Charles Sanders Peirce.

12.2 Testing for validity and entailment

When we use truth tables to test for validity or entailment, we are checking for *bad* lines: lines where the premises are all true and the conclusion is false. Note:

- Any line where the conclusion is true is not a bad line.
- Any line where some premise is false is not a bad line.

Since *all* we are doing is looking for bad lines, we should bear this in mind. So: if we find a line where the conclusion is true, we do not need to evaluate anything else on that line: that line definitely isn't bad. Likewise, if we find a line where some premise is false, we do not need to evaluate anything else on that line.

With this in mind, consider how we might test the following for validity:

$$\neg L \rightarrow (J \vee L), \neg L \therefore J$$

The *first* thing we should do is evaluate the conclusion. If we find that the conclusion is *true* on some line, then that is not a bad line. So we can simply ignore the rest of the line. So at our first stage, we are left with something like:

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

where the blanks indicate that we are not going to bother doing any more investigation (since the line is not bad) and the question-marks indicate that we need to keep investigating.

The easiest premise to evaluate is the second, so we next do that:

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

Note that we no longer need to consider the third line on the table: it will not be a bad line, because (at least) one of premises is false on that line. And finally, we complete the truth table:

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

The truth table has no bad lines, so the argument is valid. (Any valuation on which all the premises are true is a valuation on which the conclusion is true.)

It might be worth illustrating the tactic again. Let us check whether the following argument is valid

$$A \vee B, \neg(A \wedge C), \neg(B \wedge \neg D) \therefore (\neg C \vee D)$$

At the first stage, we determine the truth value of the conclusion. Since this is a disjunction, it is true whenever either disjunct is true, so we can speed things along a bit. We can then ignore every line apart from the few lines where the conclusion is false.

A	B	C	D	$A \vee B$	$\neg(A \wedge C)$	$\neg(B \wedge \neg D)$	$(\neg C \vee D)$
T	T	T	T				T
T	T	T	F	?	?	?	F F
T	T	F	T				T
T	T	F	F				T T
T	F	T	T				T
T	F	T	F	?	?	?	F F
T	F	F	T				T
T	F	F	F				T T
F	T	T	T				T
F	T	T	F	?	?	?	F F
F	T	F	T				T
F	T	F	F				T T
F	F	T	T				T
F	F	T	F	?	?	?	F F
F	F	F	T				T
F	F	F	F				T T

We must now evaluate the premises. We use shortcuts where we can:

A	B	C	D	$A \vee B$	$\neg(A \wedge C)$	$\neg(B \wedge \neg D)$	$(\neg C \vee D)$
T	T	T	T				T
T	T	T	F	T	F T		F F
T	T	F	T				T
T	T	F	F				T T
T	F	T	T				T
T	F	T	F	T	F T		F F
T	F	F	T				T
T	F	F	F				T T
F	T	T	T				T
F	T	T	F	T	T F	F TT	F F
F	T	F	T				T
F	T	F	F				T T
F	F	T	T				T
F	F	T	F	F			F F
F	F	F	T				T
F	F	F	F				T T

If we had used no shortcuts, we would have had to write 256 'T's or 'F's on this table. Using shortcuts, we only had to write 37. We have saved ourselves a *lot* of work.

I have been discussing shortcuts in testing for validity. But exactly the same shortcuts can be used in testing for entailment. By employing a similar notion of *bad* lines, you can save yourself a huge amount of work.

Practice exercises

A. Using shortcuts, determine whether each sentence is a tautology, a contradiction, or neither.

1. $\neg B \wedge B$
2. $\neg D \vee D$
3. $(A \wedge B) \vee (B \wedge A)$
4. $\neg(A \rightarrow (B \rightarrow A))$
5. $A \leftrightarrow (A \rightarrow (B \wedge \neg B))$
6. $\neg(A \wedge B) \leftrightarrow A$
7. $A \rightarrow (B \vee C)$
8. $(A \wedge \neg A) \rightarrow (B \vee C)$
9. $(B \wedge D) \leftrightarrow (A \leftrightarrow (A \vee C))$

13 Partial truth tables

Sometimes, we do not need to know what happens on every line of a truth table. Sometimes, just a single line or two will do.

Tautology. In order to show that a sentence is a tautology, we need to show that it is true on every valuation. That is to say, we need to know that it comes out true on every line of the truth table. 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 false. Therefore, in order to show that some sentence is not a tautology, it is enough to provide a single valuation – a single line of the truth table – which makes the sentence false.

Suppose that we want to show that the sentence $(U \wedge T) \rightarrow (S \wedge W)$ is *not* a tautology. We set up a PARTIAL TRUTH TABLE:

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

We have only left space for one line, rather than 16, since we are only looking for one line, on which the sentence is false. For just that reason, we have filled 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 and the consequent must be false. So we fill these in on the table:

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

In order for the $(U \wedge T)$ to be true, both ‘ U ’ and ‘ T ’ must be true.

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

Now we just need to make ' $(S \wedge W)$ ' 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	$(U \wedge T) \rightarrow (S \wedge W)$
F	T	T	F	T T T F F F F

So we now have a partial truth table, which shows that ' $(U \wedge T) \rightarrow (S \wedge W)$ ' is not a tautology. Put otherwise, we have shown that there is a valuation which makes ' $(U \wedge T) \rightarrow (S \wedge W)$ ' false, namely, the valuation which makes ' S ' false, ' T ' true, ' U ' true and ' W ' false.

Contradiction. Showing that something is a contradiction requires a complete truth table: we need to show that there is no valuation which makes the sentence true; that is, we need to show that the sentence is false on every line of the truth table.

However, to show that something is *not* a contradiction, all we need to do is find a valuation which makes the sentence true, and a single line of a truth table will suffice. We can illustrate this with the same example.

S	T	U	W	$(U \wedge T) \rightarrow (S \wedge W)$
				T

To make the sentence true, it will suffice to ensure that the antecedent is false. Since the antecedent is a conjunction, we can just make one of them false. For no particular reason, we choose to make ' U ' false; and then we can assign whatever truth value we like to the other atomic sentences.

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

equivalence. To show that two sentences are equivalent, we must show that the sentences have the same truth value on every valuation. So this requires a complete truth table.

To show that two sentences are *not* equivalent, we only need to show that there is a valuation on which they have different truth values. So this requires only a one-line partial truth table: make the table so that one sentence is true and the other false.

Consistency. To show that some sentences are jointly consistent, we must show that there is a valuation which makes all of the sentence true. So this requires only a partial truth table with a single line.

To show that some sentences are jointly inconsistent, we must show that there is no valuation which makes all of the sentence true. So this requires a complete truth table: You must show that on every row of the table at least one of the sentences is false.

Validity. To show that an argument is valid, we must show that there is no valuation which makes all of the premises true and the conclusion false. So this requires a complete truth table. (Likewise for entailment.)

To show that argument is *invalid*, we must show that there is a valuation which makes all of the premises true and the conclusion false. So this requires only a one-line partial truth table on which all of the premises are true and the conclusion is false. (Likewise for a failure of entailment.)

This table summarises what is required:

	Yes	No
tautology?	complete truth table	one-line partial truth table
contradiction?	complete truth table	one-line partial 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
entailment?	complete truth table	one-line partial truth table

Practice exercises

A. Use complete or partial truth tables (as appropriate) to determine whether these pairs of sentences are equivalent:

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. $A \wedge \neg A, \neg B \leftrightarrow B$
6. $\neg(A \wedge B), \neg A \vee \neg B$
7. $\neg(A \rightarrow B), \neg A \rightarrow \neg B$
8. $(A \rightarrow B), (\neg B \rightarrow \neg A)$

B. Use complete or partial truth tables (as appropriate) to determine whether these sentences are jointly consistent, or jointly inconsistent:

1. $A \wedge B, C \rightarrow \neg B, C$
2. $A \rightarrow B, B \rightarrow C, A, \neg C$
3. $A \vee B, B \vee C, C \rightarrow \neg A$
4. $A, B, C, \neg D, \neg E, F$

C. Use complete or partial truth tables (as appropriate) to determine whether each argument is valid or invalid:

1. $A \vee (A \rightarrow (A \leftrightarrow A)) \therefore A$
2. $A \leftrightarrow \neg(B \leftrightarrow A) \therefore A$
3. $A \rightarrow B, B \therefore A$
4. $A \vee B, B \vee C, \neg B \therefore A \wedge C$
5. $A \leftrightarrow B, B \leftrightarrow C \therefore A \leftrightarrow C$

14 Expressiveness of Sentential

When we introduced the idea of truth-functionality in §9.1, we observed that every sentence connective in Sentential was truth-functional. As we noted, that property allows us to represent complex sentences involving only these connectives using truth tables.

Are there other truth functional connectives than those in Sentential? If there were, they would have schematic truth tables that differ from those for any of our connectives. And it is easy to see that there are. Consider this proposed connective:

The Sheffer Stroke For any sentences \mathcal{A} and \mathcal{B} , ' $\mathcal{A} \downarrow \mathcal{B}$ ' is true if and only if both \mathcal{A} and \mathcal{B} are false. We can summarize this in the schematic truth table for the Sheffer Stroke:

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

Inspection of the schematic truth tables for \wedge , \vee , etc., shows that their truth tables are different from this one, and hence the Sheffer Stroke is not one of the connectives of Sentential. It is a connective of English however: it is the 'neither ... nor ...' connective that features in 'Siya is neither an archer nor a jockey', which is false iff she is either.

'Whether or not' The connective '... whether or not ...', as in the sentence 'Sam is happy whether or not she's rich' seems to have this schematic truth table:

\mathcal{A}	\mathcal{B}	' \mathcal{A} whether or not \mathcal{B} '
T	T	T
T	F	T
F	T	F
F	F	F

This too corresponds to no existing connective of Sentential.

In fact, it can be shown that there are many other potential truth-functional connectives that are not included in the language Sentential.¹

Should we be worried about this, and attempt to add new connectives to Sentential? It turns out we already have enough connectives to say *anything* we want to say using a truth-functional connective. The connectives already in Sentential are TRUTH-FUNCTIONALLY COMPLETE:

For any truth-functional connective in any language – that is, which has a truth-table – there is a schematic sentence of Sentential which has the same truth table.

Remember that a schematic sentence is something like ' $\mathcal{A} \vee \neg\mathcal{B}$ ', where arbitrary Sentential sentences can fill the places indicated by \mathcal{A} and \mathcal{B} .

This is actually not very difficult to show. We'll start with an example. Suppose we have the English sentence 'Siya is exactly one of an archer and a jockey'. This sentence features the connective 'Exactly one of ... and ...'. In this case, the simpler sentences which are connected to form the complex sentence are 'Siya is an archer' and 'Siya is a jockey'. The schematic truth table for this connective is as follows:

\mathcal{A}	\mathcal{B}	'Exactly one of \mathcal{A} and \mathcal{B} '
T	T	F
T	F	T
F	T	T
F	F	F

We want now to find a schematic Sentential sentence that has this same truth table. So we shall want the sentence to be true on the second row, and true on the third row, and false on the other rows. In other words, we want a sentence which is true on *either* the second row *or* the third row.

Let's begin by focusing on that second row, or rather the family of valuations corresponding to it. Those valuations include only those that make \mathcal{A} true and \mathcal{B} false. These are the only valuations among those we are considering which make \mathcal{A} true and \mathcal{B} false. So they are the only valuations which make both \mathcal{A} and $\neg\mathcal{B}$ true. So we can construct a sentence which is true on valuations in that family, and those valuations alone: the conjunction of \mathcal{A} and $\neg\mathcal{B}$, $(\mathcal{A} \wedge \neg\mathcal{B})$.

¹ There are in fact sixteen truth-functional connectives that join two simpler sentences into a complex sentence, but Sentential only includes four. (Why sixteen? Because there are four rows of the schematic truth table for such a connective, and each row can have either a T or an F recorded against it, independent of the other rows, so there are $2 \times 2 \times 2 \times 2 = 16$ ways of constructing such a truth-table.)

Now look at the third row and its associated family of valuations. Those valuations make \mathcal{A} false and \mathcal{B} true. They are the only valuations among those we are considering which make \mathcal{A} false and \mathcal{B} true. So they are the only valuations among those we are considering which make both of $\neg\mathcal{A}$ and \mathcal{B} true. So we can construct a schematic sentence which is true on valuations in that family, and only those valuations: the conjunction of $\neg\mathcal{A}$ and \mathcal{B} , $(\neg\mathcal{A} \wedge \mathcal{B})$.

Our target sentence, the one with the same truth table as ‘Exactly one of \mathcal{A} and \mathcal{B} ’, is true on either the second or third valuations. So it is true if either $(\mathcal{A} \wedge \neg\mathcal{B})$ is true or if $(\neg\mathcal{A} \wedge \mathcal{B})$ is true. And there is of course a schematic Sentential sentence with just this profile: $(\mathcal{A} \wedge \neg\mathcal{B}) \vee (\neg\mathcal{A} \wedge \mathcal{B})$.

Let us summarise this construction by adding to our truth table:

\mathcal{A}	\mathcal{B}	‘Exactly one of \mathcal{A} and \mathcal{B} ’	$(\mathcal{A} \wedge \neg\mathcal{B}) \vee (\neg\mathcal{A} \wedge \mathcal{B})$
T	T	F	F F F
T	F	T	T T F
F	T	T	F T T
F	F	F	F F F

As we can see, we have come up with a schematic Sentential sentence with the intended truth table.

This procedure generalises:

1. First, identify the truth table of the target connective;
2. Then, identify which families of valuations (schematic truth table rows) the target sentence is true on, and for each such row, construct a conjunctive schematic Sentential sentence true on that row alone. (It will be a conjunction of schematic letters sentences of those schematic letters which are true on the valuation, and negated schematic letters, for those schematic letters false on the valuation).
 - › What if the target connective is true on no valuations? Then let the schematic Sentential sentence $(\mathcal{A} \wedge \neg\mathcal{A})$ represent it – it too is true on no valuation.
3. Finally, the schematic Sentential sentence will be a disjunction of those conjunctions, because the target sentence is true according to any of those valuations.
 - › What if there is only one such conjunction, because the target sentence is true in only one valuation? Then just take that conjunction to be the Sentential rendering of the target sentence.

Logicians say that the schematic sentences that this procedure spits out are in DISJUNCTIVE NORMAL FORM.

This procedure doesn't always give the simplest schematic Sentential sentence with a given truth table, but for any truth table you like this procedure gives us a schematic Sentential sentence with that truth table. Indeed, we can see that the Sentential sentence $\neg(\mathcal{A} \leftrightarrow \mathcal{B})$ has the same truth table as our target sentence too.

The procedure can be used to show that there is some redundancy in Sentential itself. Take the connective \leftrightarrow . Our procedure, applied to the schematic truth table for $\mathcal{A} \leftrightarrow \mathcal{B}$, yields the following schematic sentence:

$$(\mathcal{A} \wedge \mathcal{B}) \vee (\neg\mathcal{A} \wedge \neg\mathcal{B}).$$

This schematic sentence says the same thing as the original schematic sentence with the biconditional as its main connective, without using the biconditional. This could be used as the basis of a program to remove the biconditional from the language. But that would make Sentential more difficult to use, and we will not pursue this idea further.

Practice exercises

A. For each of columns (i), (ii) and (iii) below, use the procedure just outlined to find a Sentential sentence that has the truth table depicted:

\mathcal{A}	\mathcal{B}	(i)	(ii)	(iii)
T	T	T	T	F
T	F	T	T	T
F	T	T	F	F
F	F	T	T	T

B. Can you find Sentential schematic sentences which have the same truth table as these English connectives?

1. ' \mathcal{A} whether or not \mathcal{B} ';
2. 'Not both \mathcal{A} and \mathcal{B} ';
3. 'Neither \mathcal{A} nor \mathcal{B} , but at least one of them'.

Chapter 4
Quantified logic

15 Building blocks of Quantifier

15.1 The need to decompose sentences

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

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

To symbolise it in `Sentential`, we might offer a symbolisation key:

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

And the argument itself becomes:

$$L, A \therefore F$$

This is *invalid* in `Sentential`. But the original English argument is clearly valid.

The problem is not that we have made a mistake while symbolising the argument. This is the best symbolisation we can give in `Sentential`. The problem lies with `Sentential` itself. 'All logicians wear funny hats' is about both logicians and hat-wearing. By not retaining this structure in our symbolisation, we lose the connection between Willard's being a logician and Willard's wearing a hat.

The basic units of `Sentential` are atomic sentences, and `Sentential` cannot decompose these. To symbolise arguments like the preceding one, we will have to develop a new logical language which will allow us to *split the atom*. We will call this language `Quantifier`, and the study of this language and its features is *quantified logic*.

The details of `Quantifier` will be explained throughout this chapter, but here is the basic idea for splitting the atom.

First, we have *names*. In `Quantifier`, we indicate these with lower case italic letters. For instance, we might let '*b*' stand for Bertie, or let '*i*' stand for Willard. The names of `Quantifier` correspond to proper names in English.

Second, we have predicates. English predicates are expressions like ‘_____ is a dog’ or ‘_____ is a logician’. These are not complete sentences by themselves. In order to make a complete sentence, we need to fill in the gap. We need to say something like ‘Bertie is a dog’ or ‘Willard is a logician’. In *Quantifier*, we indicate predicates with upper case italic letters. For instance, we might let the *Quantifier* predicate ‘*D*’ symbolise the English predicate ‘_____ is a dog’. Then the expression ‘*Db*’ will be a sentence in *Quantifier*, which symbolises the English sentence ‘Bertie is a dog’. Equally, we might let the *Quantifier* predicate ‘*L*’ symbolise the English predicate ‘_____ is a logician’. Then the expression ‘*Li*’ will symbolise the English sentence ‘Willard is a logician’.

Third, we have quantifiers. For instance, ‘ \exists ’ will roughly convey ‘There is at least one ...’. So we might symbolise the English sentence ‘there is a dog’ with the *Quantifier* sentence ‘ $\exists xDx$ ’, which we would naturally read out-loud as ‘there is at least one thing, x , such that x is a dog’.

That is the general idea. But *Quantifier* is significantly more subtle than *Sentential*. So we will come at it slowly.

15.2 Names

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 ‘Bertie’ is a singular term, because it refers to a specific terrier. Likewise, the phrase ‘Philip’s dog Bertie’ is a singular term, because it refers to a specific little terrier.

Proper names are a particularly important kind of singular term. These are expressions that pick out individuals without describing them. 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 ‘Hilary’ is used as a singular term, you might guess that it refers to a woman. You might, though, be guessing wrongly. Indeed, the name does not necessarily mean that the person referred to is even a person: Hilary might be a giraffe, for all you could tell just from the name.

In *Quantifier*, our NAMES are lower-case letters ‘*a*’ through to ‘*r*’. We can add subscripts if we want to use some letter more than once. So here are some singular terms in *Quantifier*:

$$a, b, c, \dots, r, a_1, f_{32}, j_{390}, m_{12}.$$

These should be thought of along the lines of proper names in English. But with one difference. ‘Antony Eagle’ is a proper name, but there are a number of people with this name. (Equally, there are at least two people with the name ‘P.D. Magnus’ and several people named ‘Tim Button’.) We live with this kind of ambiguity in English, allowing context to individuate the fact that ‘Antony Eagle’ refers to one of the contributors to this book, and not some other guy. In *Quantifier*, we do not tolerate any such ambiguity. Each name must pick out *exactly* one thing. (However, two different names may pick out the same thing.)

As with Sentential, we can provide symbolisation keys. These indicate, temporarily, what a name shall pick out. So we might offer:

e: Elsa
g: Gregor
m: Marybeth

You may have been taught in school that a noun is a naming word. It is safe to use names in Quantifier to symbolise proper nouns. But dealing with common nouns is a more subtle matter. Even though they are often described as ‘general names’, common nouns actually function as names relatively rarely. One type of common noun which often functions as a name are NATURAL KIND TERMS like ‘gold’ and ‘tiger’. These are common nouns which really do name a kind of thing. Consider this example:

Gold is scarce;
 Nothing scarce is cheap;
 So: Gold isn’t cheap.

This argument is valid in virtue of its form. The form of the first premise has the phrase ‘_____ is scarce’ being predicated of ‘gold’, in which case, ‘gold’ must be functioning as a name in this argument. But ‘_____ is gold’ is a perfectly good predicate too, so we cannot simply treat all natural kind terms as proper names of general kinds.

15.3 Predicates

The simplest predicates denote properties of individuals. They are things you can say about an object. Here are some examples of English predicates:

_____ is a dog
 _____ is a member of Monty Python
 A piano fell on _____

In general, you can think about predicates as things which combine with singular terms to make sentences. 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 obtain any of three different predicates:

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

Quantifier predicates are capital letters A through Z , with or without subscripts. We might write a symbolisation key for predicates thus:

Ax : _____ x is angry
 Hx : _____ x is happy

(Why the subscripts on the gaps? We shall return to this in §17.)

If we combine our two symbolisation keys, we can start to symbolise some English sentences that use these names and predicates in combination. For example, consider the English sentences:

1. Elsa is angry.
2. Gregor and Marybeth are angry.
3. If Elsa is angry, then so are Gregor and Marybeth.

Sentence **1** is straightforward: we symbolise it by ' Ae '.

Sentence **2**: this is a conjunction of two simpler sentences. The simple sentences can be symbolised just by ' Ag ' and ' Am '. Then we help ourselves to our resources from Sentential, and symbolise the entire sentence by ' $Ag \wedge Am$ '. This illustrates an important point: Quantifier has all of the truth-functional connectives of Sentential.

Sentence **3**: this is a conditional, whose antecedent is sentence **1** and whose consequent is sentence **2**. So we can symbolise this with ' $Ae \rightarrow (Ag \wedge Am)$ '.

Actually we can make an elegant further assumption to incorporate Sentential entirely within Quantifier. A predicate combines with singular terms to make sentences. But some English predicates combine with *zero* singular terms to make a sentence. We see this in weather expressions like 'it is snowing', in which the pronoun 'it' doesn't actually refer to anything. This predicate does not denote a property of individuals, because there is no way to attach it to a particular individual. (Try substituting some proper name for 'it' in 'it is raining' and see what nonsense you get.) Such predicates rather are already sentences by themselves, symbolised by capital letters ' A ' through ' Z ' without any adjacent name. This is exactly how we treated the grammar of atomic sentences of Sentential. So we can in fact simply include all atomic sentences of Sentential among these special predicates of Quantifier. Since Quantifier has all the connectives of Sentential and all the atomic sentences, we see that *every sentence of Sentential is also a sentence of Quantifier*.

15.4 Quantifiers

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

4. Everyone is happy.
5. Someone is angry.

It might be tempting to symbolise sentence 4 as $He \wedge Hg \wedge Hm$. Yet this would only say that Elsa, Gregor, and Marybeth are happy. We want to say that *everyone* is happy, even those with no names.

Note that 4 and 5 can be roughly paraphrased like this:

6. Every person is such that: they are happy.
7. Some person is such that: they are angry.

In each of these, we have an instance of singular ‘they’ which is governed by the preceding phrase. That phrase tells us who this instance of ‘they’ is pointing to – is it pointing severally to everyone, as in 6? Or to just someone, though perhaps it is unknown which particular person it is, as in 7? In this sort of construction, the sentence ‘they are happy’ headed by the bare pronoun is called an OPEN SENTENCE. An open sentence in English can be used to say something meaningful, if the circumstances permit a unique interpretation of the pronoun. But in many cases no such unique interpretation is possible. If I gesture at a crowd and say simply ‘he is angry’, I may not manage to say anything meaningful if there is no way to establish which person this use of ‘he’ is pointing to. The other part of the sentence, the ‘every person is such that ...’ part, is called a QUANTIFIER PHRASE. The occurrence of the quantifier phrase gives us guidance about how to interpret the pronoun.

The treatment of quantifier phrases in Quantifier actually follows the structure of these paraphrases rather well. The role of the pronoun is played by a VARIABLE. In Quantifier, variables are italic lower case letters ‘s’ through ‘z’, with or without subscripts. They combine with predicates to form open sentences of the form $\mathcal{A}x$. An open sentence combines with a quantifier phrase to form a sentence. (Notice that I have here returned to the practice of using \mathcal{A} as a metavariable, from §7.)

The first quantifier from Quantifier we meet is the UNIVERSAL QUANTIFIER, symbolised ‘ \forall ’, and which corresponds to ‘every’. Unlike English, we always follow a quantifier in Quantifier by the variable it governs, to avoid the possibility of confusion. Putting this all together, we might symbolise sentence 4 as $\forall xHx$. The variable ‘ x ’ is serving as a kind of placeholder. The expression ‘ $\forall x$ ’ intuitively means that you can pick anyone to be temporarily denoted by ‘ x ’. The subsequent ‘ Hx ’ indicates, of that thing you picked out, that *it* is happy. (Note that pronoun again.)

I should say that there is no special reason to use ‘ x ’ rather than some other variable. The sentences ‘ $\forall xHx$ ’, ‘ $\forall yHy$ ’, ‘ $\forall zHz$ ’, and ‘ $\forall x_5Hx_5$ ’ use different variables, but they will all be logically equivalent.

To symbolise sentence 5, we introduce a second quantifier: the EXISTENTIAL QUANTIFIER, ‘ \exists ’. Like the universal quantifier, the existential quantifier requires a variable. Sentence 5 can be symbolised by ‘ $\exists xAx$ ’. Whereas ‘ $\forall xAx$ ’ is read naturally as ‘for all x , it (x) is angry’, ‘ $\exists xAx$ ’ is read naturally as ‘there is something, x , such that it (x) is angry’. Once again, the variable is a kind of placeholder; we could just as easily have symbolised sentence 5 with ‘ $\exists zAz$ ’, ‘ $\exists w_{256}Aw_{256}$ ’, or whatever.

Some more examples will help. Consider these further sentences:

8. No one is angry.
9. There is someone who is not happy.
10. Not everyone is happy.

Sentence 8 can be paraphrased as, ‘It is not the case that someone is angry’. We can then symbolise it using negation and an existential quantifier: $\neg\exists xAx$. Yet sentence 8 could also be paraphrased as, ‘Everyone is not angry’. With this in mind, it can be symbolised using negation and a universal quantifier: $\forall x\neg Ax$. Both of these are acceptable symbolisations. Indeed, it will transpire that, in general, $\forall x\neg A$ is logically equivalent to $\neg\exists xA$. Symbolising a sentence one way, rather than the other, might seem more ‘natural’ in some contexts, but it is not much more than a matter of taste.

Sentence 9 is most naturally paraphrased as, ‘There is some x , such that x is not happy’. This then becomes $\exists x\neg Hx$. Of course, we could equally have written $\neg\forall xHx$, which we would naturally read as ‘it is not the case that everyone is happy’. And that would be a perfectly adequate symbolisation of sentence 10.

Quantifiers get their name because they tell us how many things have a certain feature. Quantifier allows only very crude distinctions: we have seen that we can symbolise ‘no one’, ‘someone’, and ‘everyone’. English has many other quantifier phrases: ‘most’, ‘a few’, ‘more than half’, ‘at least three’, etc. Some can be handled in a roundabout way in Quantifier, as we will see: the numerical quantifier ‘at least three’, for example, we will meet again in §18. But others, like ‘most’, are simply unable to be reliably symbolised in Quantifier.

15.5 Domains

Given the symbolisation key we have been using, $\forall xHx$ symbolises ‘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 usually mean something more modest: everyone now in the building, everyone enrolled in the ballet class, or whatever.

In order to eliminate this ambiguity, we will need to specify a DOMAIN. The domain is just the things that we are talking about. So if we want to talk about people in Chicago, we define the domain to be people in Chicago. We write this at the beginning of the symbolisation key, like this:

domain: people in Chicago

The quantifiers *range over* the domain. Given this domain, $\forall x$ is to be read roughly as ‘Every person in Chicago is such that ...’ and $\exists x$ is to be read roughly as ‘Some person in Chicago is such that ...’

In Quantifier, the domain must always include at least one thing. Moreover, in English we can conclude ‘something is angry’ when given ‘Gregor is angry’. In Quantifier, then, we shall want to be able to infer ‘ $\exists xAx$ ’ from ‘ Ag ’. So we shall insist that each name must pick out exactly one thing in the domain. If we want to name people in places beside Chicago, then we need to include those people in the domain.

A domain must have *at least* one member. A name must pick out *exactly* one member of the domain. But a member of the domain may be picked out by one name, many names, or none at all.

In permitting multiple domains, Quantifier follows the lead of natural languages like English. Consider an argument like this:

11. All the beer has been drunk; so we’re going to the bottle-o.

The premise says that all the beer is gone. But the conclusion only makes sense if there is more beer at the bottle shop. So whatever domain of things we are talking about when we state the premise, it cannot include absolutely *everything*. In Quantifier, we sidestep the interesting issues involved in deciding just what domain is involved in evaluating sentences like ‘all the beer has been drunk’, and explicitly include the domain of quantification in our symbolisation key.

16 Sentences with one quantifier

We now have all of the pieces of Quantifier. Symbolising more complicated sentences will only be a matter of knowing the right way to combine predicates, names, quantifiers, and connectives. There is a knack to this, and there is no substitute for practice.

16.1 Dealing with complex adjectives

When we encounter a sentence like

1. Herbie is a white car,

we can paraphrase this as ‘Herbie is white and Herbie is a car’. We can then use a symbolisation key like:

Wx : _____ x is white
 Cx : _____ x is a car
 h : Herbie

This allows us to symbolise sentence 1 as ‘ $Wh \wedge Ch$ ’. But now consider:

2. Julia Gillard is a former prime minister.
3. Julia Gillard is prime minister.

Following the case of Herbie, we might try to use a symbolisation key like:

Fx : _____ x is former
 Px : _____ x is Prime Minister
 j : Julia Gillard.

Then we would symbolise 2 by ‘ $Fj \wedge Pj$ ’, and symbolise 3 by ‘ Pj ’. That would however be a mistake, since that symbolisation suggests that the argument from 2 to 3 is valid, because the symbolisation of the premise does logically entail the symbolisation of the conclusion.

‘White’ is a SUBJECTIVE adjective, which is a fancy way of saying that the white F s are among the F s: any white car is a car, just like any successful lawyer is a lawyer. But ‘former’

is a PRIVATIVE adjective, which means that any former F is not now among the F s. Other privative adjectives occur in phrases such as ‘fake diamond’, ‘Deputy Lord Mayor’, and ‘mock trial’. When symbolising these sentences, you cannot treat them as a conjunction. So you will need to symbolise ‘_____ is a fake diamond’ and ‘_____ is a diamond’ using completely different predicates, to avoid a spurious entailment between them. The moral is: when you see two adjectives in a row, you need to ask yourself carefully whether they are subsective, and can be treated as a conjunction, or not.¹

Things are a bit more complicated, however. Recall this example from page 56:

Daisy is a small cow.

We note that a small cow is definitely a cow, and so it seems we can treat ‘small’ as a subsective adjective. We might formalise this sentence like this: ‘ $Sd \wedge Cd$ ’, assuming this symbolisation key:

Sx : _____ x is small
 Cx : _____ x is a cow
 d : Daisy

But note that our symbolisation would suggest that this argument is valid:

4. Daisy is a small cow; so Daisy is small.

The symbolised argument, $Sd \wedge Cd \therefore Sd$, is clearly valid.

But the original argument 4 is not valid. Even a small cow is still rather large. (Likewise, even a short basketball player is still generally well above average height.) The point is that ‘_____ x is a small cow’ denotes the property something has when it is small *for a cow*, while ‘_____ x is small’ denotes the property of being a small *thing*. (In ordinary speech we tend to keep the ‘for an F ’ part of these phrases silent, and let our conversational circumstances supply it automatically.) But neither should we treat ‘small’ as a non-subsective adjective. If we do, we will be unable to account for the valid argument ‘Daisy is a small cow, so Daisy is a cow’.

The correct symbolisation key will thus be this, keeping the other symbols as they were:

Sx : _____ x is small-for-a-cow

On this symbolisation key, the valid formal argument $Sd \wedge Cd \therefore Cd$ corresponds to this valid argument, as it should:

¹ This caution also applies to adjectives which are neither subsective nor privative, like ‘alleged’ in ‘alleged murderer’. These ought not be symbolised by conjunction either.

Daisy is a cow that is small-for-a-cow; so Daisy is a cow.

Likewise, this rather unusual English argument turns out to be valid too:

Daisy is a cow that is small-for-a-cow; so Daisy is small-for-a-cow.

(Note that it can be rather difficult to hear the English sentence ‘Daisy is small’ as saying the same thing as the conclusion of this argument, ‘Daisy is small for a cow’, which explains why ‘Daisy is a small cow, so Daisy is small’ strikes us as invalid.)

If we take these observations to heart, there are many substantive adjectives which can change their meaning depending on what predicate they are paired with. Small-for-an-oil-tanker is a rather different size property than small-for-a-mouse, but in ordinary English we use the phrases ‘small oil tanker’ and ‘small mouse’ without bothering to make these different senses of ‘small’ explicit. When symbolising, it is best to make them very explicit, because doing so blocks the fallacious argument from ‘Daisy is a small cow’ to ‘Daisy is small’, where the natural sense of the conclusion is the generic size claim ‘Daisy is small-for-a-thing’.

The upshot is this: though you can symbolise ‘Daisy is a small cow’ as a conjunction, you will need to symbolise ‘_____ is a small cow’ and ‘_____ is a small animal’ using different predicates of Quantifier to stand for the different appearances of ‘small’ – to symbolise ‘small for a cow’ and ‘small for an animal’.

The overall message is not particularly specific: treat adjectives with care, and always think about whether a conjunction or some other symbolisation best captures what is going on in the English. There is no substitute for practice in developing a good sense of how to symbolise arguments.

16.2 Common quantifier phrases

Consider these sentences:

5. Every coin in my pocket is a quarter.
6. Some coin on the table is a dime.
7. Not all the coins on the table are dimes.
8. None of the coins in my pocket are dimes.

In providing a symbolisation key, we need to specify a domain. Since we are talking about coins in my pocket and on the table, the domain must at least contain all of those coins. Since we are not talking about anything besides coins, we let the domain be all coins. Since we are not talking about any specific coins, we do not need to deal with any names. So here is our key:

domain: 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 5 is most naturally symbolised using a universal quantifier. The universal quantifier says something about everything in the domain, not just about the coins in my pocket. Sentence 5 can be paraphrased as ‘for any coin, *if* that coin is in my pocket *then* it is a quarter’. So we can symbolise it as ‘ $\forall x(Px \rightarrow Qx)$ ’.

Since sentence 5 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 ‘ $\forall x(Px \wedge Qx)$ ’ would symbolise the sentence ‘every coin is both a quarter and in my pocket’. This obviously means something very different than sentence 5. And so we see:

A sentence can be symbolised as $\forall x(Fx \rightarrow Gx)$ if it can be paraphrased in English as ‘every F is G’.

Sentence 6 is most naturally symbolised using an existential quantifier. It can be paraphrased as ‘there is some coin which is both on the table and which is a dime’. So we can symbolise it as ‘ $\exists x(Tx \wedge Dx)$ ’.

Notice that we needed to use a conditional with the universal quantifier, but we used a conjunction with the existential quantifier. Suppose we had instead written ‘ $\exists x(Tx \rightarrow Dx)$ ’. That would mean that there is some object in the domain such that *if* it is T , then it is also D . For this to be true, we just need something to not be T . So it is *very easy* for ‘ $\exists x(Tx \rightarrow Dx)$ ’ to be true. Given our symbolisation, it will be true if some coin is not on the table. Of course there is a coin that is not on the table: there are coins lots of other places. That is rather less demanding than the claim that something is both T and D .

A conditional will usually be the natural connective to use with a universal quantifier, but a conditional within the scope of an existential quantifier tends to say something very weak indeed. As a general rule of thumb, do not put conditionals in the scope of existential quantifiers unless you are sure that you need one.

A sentence can be symbolised as $\exists x(Fx \wedge Gx)$ if it can be paraphrased in English as ‘some F is G’.

Sentence 7 can be paraphrased as, ‘It is not the case that every coin on the table is a dime’. So we can symbolise it by ‘ $\neg \forall x(Tx \rightarrow Dx)$ ’. You might look at sentence 7 and paraphrase it instead as, ‘Some coin on the table is not a dime’. You would then symbolise it by ‘ $\exists x(Tx \wedge \neg Dx)$ ’. Although it is probably not immediately obvious yet, these two sentences are logically equivalent. (This is due to the logical equivalence between $\neg \forall x \mathcal{A}$ and $\exists x \neg \mathcal{A}$, mentioned in §15, along with the equivalence between $\neg(\mathcal{A} \rightarrow \mathcal{B})$ and $\mathcal{A} \wedge \neg \mathcal{B}$.)

Sentence 8 can be paraphrased as, ‘It is not the case that there is some dime in my pocket’. This can be symbolised by ‘ $\neg\exists x(Px \wedge Dx)$ ’. It might also be paraphrased as, ‘Everything in my pocket is a non-dime’, and then could be symbolised by ‘ $\forall x(Px \rightarrow \neg Dx)$ ’. Again the two symbolisations are logically equivalent. Both are correct symbolisations of sentence 8.

16.3 Empty predicates

In §15, I emphasised that a name must pick out exactly one object in the domain. However, a predicate need not apply to anything in the domain. A predicate that applies to nothing in the domain is called an EMPTY predicate. This is worth exploring.

Suppose we want to symbolise these two sentences:

9. Every monkey knows sign language
10. Some monkey knows sign language

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

domain: animals

Mx : _____ x is a monkey.

Sx : _____ x knows sign language.

Sentence 9 can now be symbolised by ‘ $\forall x(Mx \rightarrow Sx)$ ’. Sentence 10 can be symbolised as ‘ $\exists x(Mx \wedge Sx)$ ’.

It is tempting to say that sentence 9 *entails* sentence 10. That is, we might think that it is impossible for it to be the case that every monkey knows sign language, without it’s also being the case that some monkey knows sign language. But this would be a mistake. It is possible for the sentence ‘ $\forall x(Mx \rightarrow Sx)$ ’ to be true even though the sentence ‘ $\exists x(Mx \wedge Sx)$ ’ is false.

How can this be? The answer comes from considering whether these sentences would be true or false *if there were no monkeys*. If there were no monkeys at all (in the domain), then ‘ $\forall x(Mx \rightarrow Sx)$ ’ would be *vacuously* true: take any monkey you like – it knows sign language! But if there were no monkeys at all (in the domain), then ‘ $\exists x(Mx \wedge Sx)$ ’ would be false.

In *Quantifier*, a universally quantified sentence of the form $\forall xAx \rightarrow Bx$ is *false* only if we can find something which is A without being B . If we can’t find such a thing, perhaps because we can’t find anything which is A in the first place, then the sentence will be true (since truth is just lack of falsity, and this sentence isn’t false because we can’t find a case that falsifies it).

Another example will help to bring this home. Suppose we extend the above symbolisation key, by adding:

Rx : _____ x is a refrigerator

Now consider the sentence ' $\forall x(Rx \rightarrow Mx)$ '. This symbolises 'every refrigerator is a monkey'. And this sentence is true, given our symbolisation key. 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 ' $\forall x(Rx \rightarrow Mx)$ ' is true iff any member of the domain that is a refrigerator is a monkey. Since the domain is *animals*, there are no refrigerators in the domain. Again, then, the sentence is *vacuously* true.

If you were actually dealing with the sentence 'All refrigerators are monkeys', then you would most likely want to include kitchen appliances in the domain. Then the predicate ' R ' would not be empty and the sentence ' $\forall x(Rx \rightarrow Mx)$ ' would be false. Remember, though, that a predicate is empty only *relative* to a particular domain.

When \mathcal{F} is an empty predicate relative to a given domain, a sentence $\forall x(\mathcal{F}x \rightarrow \dots)$ will be vacuously true of that domain.

16.4 Picking a domain

The appropriate symbolisation of an English language sentence in *Quantifier* will depend on the symbolisation key. Choosing a key can be difficult. Suppose we want to symbolise the English sentence:

- ii. Every rose has a thorn.

We might offer this symbolisation key:

Rx : _____ x is a rose
 Tx : _____ x has a thorn

It is tempting to say that sentence ii should be symbolised as ' $\forall x(Rx \rightarrow Tx)$ '. But we have not yet chosen a domain. If the domain contains all roses, this would be a good symbolisation. Yet if the domain is merely *things on my kitchen table*, then ' $\forall x(Rx \rightarrow Tx)$ ' would only come close to covering the fact 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. This is not what we want. To symbolise sentence ii adequately, we need to include all the roses in the domain. But now we have two options.

First, we can restrict the domain to include all roses but *only* roses. Then sentence ii can, if we like, be symbolised with ' $\forall xTx$ '. This is true iff everything in the domain has a thorn; since the domain is just the roses, this is true iff every rose has a thorn. By restricting the domain, we have been able to symbolise our English sentence with a very short sentence of *Quantifier*. So this approach can save us trouble, if every sentence that we want to deal with is about roses.

Second, we can let the domain contain things besides roses: rhododendrons; rats; rifles; whatever. And we will certainly need to include a more expansive domain if we simultaneously want to symbolise sentences like:

12. Every cowboy sings a sad, sad song.

Our domain must now include both all the roses (so that we can symbolise sentence 11) and all the cowboys (so that we can symbolise sentence 12). So we might offer the following symbolisation key:

domain: people and plants

Cx : _____ x is a cowboy

Sx : _____ x sings a sad, sad song

Rx : _____ x is a rose

Tx : _____ x has a thorn

Now we will have to symbolise sentence 11 with ' $\forall x(Rx \rightarrow Tx)$ ', since ' $\forall xTx$ ' would symbolise the sentence 'every person or plant has a thorn'. Similarly, we will have to symbolise sentence 12 with ' $\forall x(Cx \rightarrow Sx)$ '.

In general, the universal quantifier can be used to symbolise the English expression 'everyone' if the domain only contains people. If there are people and other things in the domain, then 'everyone' must be treated as 'every person'.

16.5 The utility of paraphrase

When symbolising English sentences in *Quantifier*, it is important to understand the structure of the sentences you want to symbolise. What matters is the final symbolisation in *Quantifier*, and sometimes you will be able to move from an English language sentence directly to a sentence of *Quantifier*. 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 finally symbolise directly in *Quantifier*.

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

domain: people

Bx : _____ x is a bassist.

Rx : _____ x is a rock star.

k : Kim Deal

Now consider these sentences:

13. If Kim Deal is a bassist, then she is a rock star.
14. If a woman is a bassist, then she is a rock star.

The same words appear as the consequent in sentences 13 and 14 ('... she is a rock star'), but they mean very different things. To make this clear, it often helps to paraphrase the original sentences into a more unusual but clearer form.

Sentence 13 can be paraphrased as, 'Consider Kim Deal: if she is a bassist, then she is a rockstar'. The bare pronoun 'she' gets to denote Kim Deal because of our initial 'Consider Kim Deal' remark. This then says something about one particular person, and can obviously be symbolised as ' $Bk \rightarrow Rk$ '.

Sentence 14 gets a very similar paraphrase, with the same embedded conditional: 'Consider any woman: if she is a bassist, then she is a rockstar'. The difference in the 'Consider ...' phrase however forces a very different interpretation for the sentence as a whole. Replacing the English pronouns by variables, the Quantifier equivalent of a pronoun, we get this awkward quasi-English paraphrase: 'For any woman x , if x is a bassist, then x is a rockstar'. Now this can be symbolised as ' $\forall x(Bx \rightarrow Rx)$ '. This is the same sentence we would have used to symbolise 'Every woman who is a bassist is a rock star'. And on reflection, that is surely true iff sentence 14 is true, as we would hope.

Consider these further sentences:

15. If anyone is a bassist, then Kim Deal is a rock star.
16. If anyone is a bassist, then they are a rock star.

The same words appear as the antecedent in sentences 15 and 16 ('If anyone is a bassist...'). But it can be tricky to work out how to symbolise these two uses. Again, paraphrase will come to our aid.

Sentence 15 can be paraphrased, 'If there is at least one bassist, then Kim Deal is a rock star'. It is now clear that this is a conditional whose antecedent is a quantified expression; so we can symbolise the entire sentence with a conditional as the main connective: ' $\exists xBx \rightarrow Rk$ '.

Sentence 16 can be paraphrased, 'For all people x , if x is a bassist, then x is a rock star'. Or, in more natural English, it can be paraphrased by 'All bassists are rock stars'. It is best symbolised as ' $\forall x(Bx \rightarrow Rx)$ ', just like sentence 14.

The moral is that the English words 'any' and 'anyone' should typically be symbolised using quantifiers. And if you are having a hard time determining whether to use an existential or a universal quantifier, try paraphrasing the sentence with an English sentence that uses words *besides* 'any' or 'anyone'.

16.6 Quantifiers and scope

Continuing the example, suppose I want to symbolise these sentences:

17. If everyone is a bassist, then Tim is a bassist
18. Everyone is such that, if they are a bassist, then Tim is a bassist.

To symbolise these sentences, I shall have to add a new name to the symbolisation key, namely:

b: Tim

Sentence 17 is a conditional, whose antecedent is ‘everyone is a bassist’. So we will symbolise it with ‘ $\forall xBx \rightarrow Bb$ ’. This sentence is *necessarily* true: if *everyone* is indeed a bassist, then take any one you like – for example Tim – and he will be a bassist.

Sentence 18, by contrast, might best be paraphrased by ‘every person *x* is such that, if *x* is a bassist, then Tim is a bassist’. This is symbolised by ‘ $\forall x(Bx \rightarrow Bb)$ ’. And this sentence is false. Kim Deal is a bassist. So ‘*Bk*’ is true. But Tim is not a bassist, so ‘*Bb*’ is false. Accordingly, ‘ $Bk \rightarrow Bb$ ’ will be false. So ‘ $\forall x(Bx \rightarrow Bb)$ ’ will be false as well.

In short, ‘ $\forall xBx \rightarrow Bb$ ’ and ‘ $\forall x(Bx \rightarrow Bb)$ ’ are very different sentences. We can explain the difference in terms of the *scope* of the quantifier. The scope of quantification is very much like the scope of negation, which we considered when discussing Sentential, and it will help to explain it in this way.

In the sentence ‘ $\neg Bk \rightarrow Bb$ ’, the scope of ‘ \neg ’ is just the antecedent of the conditional. We are saying something like: if ‘*Bk*’ is false, then ‘*Bb*’ is true. Similarly, in the sentence ‘ $\forall xBx \rightarrow Bb$ ’, the scope of ‘ $\forall x$ ’ is just the antecedent of the conditional. We are saying something like: if ‘*Bx*’ is true of *everything*, then ‘*Bb*’ is also true.

In the sentence ‘ $\neg(Bk \rightarrow Bb)$ ’, the scope of ‘ \neg ’ is the entire sentence. We are saying something like: ‘ $(Bk \rightarrow Bb)$ ’ is false. Similarly, in the sentence ‘ $\forall x(Bx \rightarrow Bb)$ ’, the scope of ‘ $\forall x$ ’ is the entire sentence. We are saying something like: ‘ $(Bx \rightarrow Bb)$ ’ is true of *everything*.

The moral of the story is simple. When you are using conditionals, be very careful to make sure that you have sorted out the scope correctly.

Practice exercises

A. Here are the syllogistic figures identified by Aristotle and his successors, along with their medieval names:

- › **Barbara**. All G are F. All H are G. So: All H are F
- › **Celarent**. No G are F. All H are G. So: No H are F
- › **Ferio**. No G are F. Some H is G. So: Some H is not F
- › **Darii**. All G are F. Some H is G. So: Some H is F.
- › **Camestres**. All F are G. No H are G. So: No H are F.
- › **Cesare**. No F are G. All H are G. So: No H are F.

- › **Baroko.** All F are G. Some H is not G. So: Some H is not F.
- › **Festino.** No F are G. Some H are G. So: Some H is not F.
- › **Datisi.** All G are F. Some G is H. So: Some H is F.
- › **Disamis.** Some G is F. All G are H. So: Some H is F.
- › **Ferison.** No G are F. Some G is H. So: Some H is not F.
- › **Bokardo.** Some G is not F. All G are H. So: Some H is not F.
- › **Camenes.** All F are G. No G are H So: No H is F.
- › **Dimaris.** Some F is G. All G are H. So: Some H is F.
- › **Fresison.** No F are G. Some G is H. So: Some H is not F.

Symbolise each argument in Quantifier.

B. Using the following symbolisation key:

domain: people

Kx : _____ x knows the combination to the safe

Sx : _____ x is a spy

Vx : _____ x is a vegetarian

h : Hofthor

i : Ingmar

symbolise the following sentences in Quantifier:

1. Neither Hofthor nor Ingmar is a vegetarian.
2. No spy knows the combination to the safe.
3. No one knows the combination to the safe unless Ingmar does.
4. Hofthor is a spy, but no vegetarian is a spy.

C. Using this symbolisation key:

domain: all animals

Ax : _____ x is an alligator.

Mx : _____ x is a monkey.

Rx : _____ x is a reptile.

Zx : _____ x lives at the zoo.

a : Amos

b : Bouncer

c : Cleo

symbolise each of the following sentences in Quantifier:

1. Amos, Bouncer, and Cleo all live at the zoo.
2. Bouncer is a reptile, but not an alligator.
3. Some reptile lives at the zoo.
4. Every alligator is a reptile.
5. Any animal that lives at the zoo is either a monkey or an alligator.
6. There are reptiles which are not alligators.
7. If any animal is a reptile, then Amos is.
8. If any animal is an alligator, then it is a reptile.

D. For each argument, write a symbolisation key and symbolise the argument in Quantifier.

1. Willard is a logician. All logicians wear funny hats. So Willard wears a funny hat
2. 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.
3. All my dreams are black and white. Old TV shows are in black and white. Therefore, some of my dreams are old TV shows.
4. 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.
5. 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.
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.

17 Multiple generality

So far, we have only considered sentences that require one-place predicates and one quantifier. The full power of Quantifier really comes out when we start to use many-place predicates and multiple quantifiers. For this insight, we largely have Gottlob Frege (1879) to thank, but also Peirce.

17.1 Many-place predicates

All of the predicates that we have considered so far concern properties that objects might have. The predicates have one gap in them, and to make a sentence, we simply need to slot in one name. They are ONE-PLACE predicates.

But other predicates concern the *relation* between two things. Here are some examples of relational predicates in English:

_____ loves _____
_____ is to the left of _____
_____ is in debt to _____

These are TWO-PLACE predicates. They need to be filled in with two terms in order to make a sentence. Conversely, if we start with an English sentence containing many singular terms, we can remove two singular terms, to obtain different two-place predicates. Consider the sentence 'Vinnie borrowed the family car from Nunzio'. By deleting two singular terms, we can obtain any of three different two-place predicates

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

And by removing all three singular terms, we obtain a THREE-PLACE predicate:

_____ borrowed _____ from _____

Indeed, there is no in principle upper limit on the number of places that our predicates may contain.

Now there is a little foible with the above. I have used the same symbol, '____', to indicate a gap formed by deleting a term from a sentence. However (as Frege emphasised), these are *different* gaps. To obtain a sentence, we can fill them in with the same term, but we can equally fill them in with different terms, and in various different orders. The following are all perfectly good sentences, and they all mean very different things:

Karl loves Karl
 Karl loves Imre
 Imre loves Karl
 Imre loves Imre

The point is that we need to keep track of the gaps in predicates, so that we can keep track of how we are filling them in.

To keep track of the gaps, I shall label them. The labelling conventions I shall adopt are best explained by example. Suppose I want to symbolise the following sentences:

1. Karl loves Imre.
2. Imre loves himself.
3. Karl loves Imre, but not vice versa.
4. Karl is loved by Imre.

I will start with the following representation key:

domain: people
i: Imre
k: Karl
 Lxy : ______{*x*} loves ______{*y*}

Sentence **1** will now be symbolised by ' Lki '.

Sentence **2** can be paraphrased as 'Imre loves Imre'. It can now be symbolised by ' Lii '.

Sentence **3** is a conjunction. We might paraphrase it as 'Karl loves Imre, and Imre does not love Karl'. It can now be symbolised by ' $Lki \wedge \neg Lik$ '.

Sentence **4** might be paraphrased by 'Imre loves Karl'. It can then be symbolised by ' Lik '. Of course, this slurs over the difference in tone between the active and passive voice; such nuances are lost in Quantifier.

This last example, though, highlights something important. Suppose we add to our symbolisation key the following:

Mxy : ______{*y*} loves ______{*x*}

Here, we have used the same English word ('loves') as we used in our symbolisation key for ' Lxy '. However, we have swapped the order of the *gaps* around (just look closely at those little subscripts!) So ' Mki ' and ' Lik ' now *both* symbolise 'Imre loves Karl'. ' Mik ' and ' Lki ' now *both* symbolise 'Karl loves Imre'. Since love can be unrequited, these are very different claims.

The moral is simple. When we are dealing with predicates with more than one place, we need to pay careful attention to the order of the places.

Suppose we add to our symbolisation key the following:

Sx : _____ x thinks only of _____ x
 Txy : _____ x thinks only of _____ y

As in the case of ' Lxy ' and ' Mxy ' above, the difference between these examples is only in how the gaps in the construction '... thinks only of ...' are labelled. In ' Txy ', we have labelled the two gaps differently. They do not *need* to be filled with different names or variables, but there is always the potential to put different names in those different gaps. In the case of ' Sx ', the gaps have the same label. In some sense, there is *only one* gap in this sentence, which is why the symbolisation key associates it with a one-place predicate – it means something like ' x thinks only of herself'.

17.2 The scope of quantifiers

Consider the sentence 'everyone loves someone'. This is potentially ambiguous. It might mean either of the following:

5. For every person x , there is some person that x loves
6. There is some particular person whom every person loves

Sentence 5 can be symbolised by ' $\forall x\exists yLxy$ ', and would be true of a love-triangle. For example, suppose that our domain of discourse is restricted to Imre, Juan and Karl. Suppose also that Karl loves Imre but not Juan, that Imre loves Juan but not Karl, and that Juan loves Karl but not Imre. Then sentence 5 is true.

Sentence 6 is symbolised by ' $\exists y\forall xLxy$ '. Sentence 6 is *not* true in the situation just described. Again, suppose that our domain of discourse is restricted to Imre, Juan and Karl. Then this requires that all of Juan, Imre and Karl converge on (at least) one object of love.

The point of the example is to illustrate that the scope of the quantifiers matters a great deal. Indeed, to switch them around is called a *quantifier shift fallacy*. Here is an example, which comes up in various forms throughout the philosophical literature:

For every person, there is some truth they cannot know. ($\forall\exists$)
 So: There is some truth, that no person can know. ($\exists\forall$)

This argument form is obviously invalid. It's just as bad as:¹

Every dog has its day.	(∀∃)
So: There is a day for all the dogs.	(∃∀)

The moral is: take great care with the scope of quantification. Later in this part, we will be a little more precise about the way that quantifier scope functions in *Quantifier*: see §20.

17.3 Stepping-stones to symbolisation

Once we have the possibility of multiple quantifiers and many-place predicates, representation in *Quantifier* can quickly start to become a bit tricky. When you are trying to symbolise a complex sentence, I recommend laying down several stepping stones. As usual, this idea is best illustrated by example. Consider this representation key:

domain: people and dogs
 Dx : _____ x is a dog
 Fxy : _____ x is a friend of _____ y
 Oxy : _____ x owns _____ y
 g : Geraldo

And now let's try to symbolise these sentences:

7. Geraldo is a dog owner.
8. Someone is a dog owner.
9. All of Geraldo's friends are dog owners.
10. Every dog owner is the friend of a dog owner.
11. Every dog owner's friend owns a dog of a friend.

Sentence 7 can be paraphrased as, 'There is a dog that Geraldo owns'. This can be symbolised by ' $\exists x(Dx \wedge Ogx)$ '.

Sentence 8 can be paraphrased as, 'There is some y such that y is a dog owner'. Dealing with part of this, we might write ' $\exists y(y \text{ is a dog owner})$ '. Now the fragment we have left as ' y is a dog owner' is much like sentence 7, except that it is not specifically about Geraldo. So we can symbolise sentence 8 by:

$$\exists y \exists x (Dx \wedge Oyx)$$

I need to pause to clarify something here. In working out how to symbolise the last sentence, we wrote down ' $\exists y(y \text{ is a dog owner})$ '. To be very clear: this is *neither* an *Quantifier* sentence *nor* an English sentence: it uses bits of *Quantifier* (' \exists ', ' y ') and bits of English

¹ Thanks to Rob Trueman for the example.

(‘dog owner’). It is really is *just a stepping-stone* on the way to symbolising the entire English sentence with a Quantifier sentence. You should regard it as a bit of rough-working-out, on a par with the doodles that you might absent-mindedly draw in the margin of this book, whilst you are concentrating fiercely on some problem.

Sentence 9 can be paraphrased as, ‘Everyone who is a friend of Geraldo is a dog owner’. Using our stepping-stone tactic, we might write

$$\forall x(Fxg \rightarrow x \text{ is a dog owner})$$

Now the fragment that we have left to deal with, ‘ x is a dog owner’, is structurally just like sentence 7. But it would be a mistake for us simply to write

$$\forall x(Fxg \rightarrow \exists x(Dx \wedge Oxx))$$

for we would here have a *clash of variables*. The scope of the universal quantifier, ‘ $\forall x$ ’, is the entire conditional, so the ‘ x ’ in ‘ Dx ’ should be governed by that. But ‘ Dx ’ also falls under the scope of the existential quantifier ‘ $\exists x$ ’, so the ‘ x ’ in ‘ Dx ’ should be governed by that. Which quantifier has priority? There is answer in Quantifier: if a variable x occurs in an Quantifier sentence, it is always the quantifier which has the narrowest scope which includes that occurrence of x that governs the variable. So in the sentence above, the quantifier ‘ $\exists x$ ’ governs each occurrence of ‘ x ’ in ‘ $(Dx \wedge Oxx)$ ’. Given this, that sentence does not mean what was intended. It says, roughly, ‘everyone who is a friend of Geraldo is such that there is a self-owning dog’. This is not at all the meaning of the English sentence we are aiming to symbolise.

To provide an adequate symbolisation, then, we must choose some different variable for our existential quantifier. What we want is something like:

$$\forall x(Fxg \rightarrow \exists z(Dz \wedge Oxz))$$

and this adequately symbolises sentence 9.

Sentence 10 can be paraphrased as ‘For any x that is a dog owner, there is a dog owner who is a friend of x ’. Using our stepping-stone tactic, this becomes

$$\forall x(x \text{ is a dog owner} \rightarrow \exists y(y \text{ is a dog owner} \wedge Fyx))$$

Completing the symbolisation, we end up with

$$\forall x(\exists z(Dz \wedge Oxz) \rightarrow \exists y(\exists z(Dz \wedge Oyz) \wedge Fyx))$$

Note that we have used the same letter, ‘ z ’, in both the antecedent and the consequent of the conditional, but that these are governed by two different quantifiers. This is ok: there is no potential confusion here, because it is obvious which quantifier governs each variable. We might graphically represent the scope of the quantifiers thus:

$$\frac{\frac{\frac{\text{scope of } \forall x'}{\frac{\text{scope of 1st } \exists z'}{\forall x(\exists z(Dz \wedge Oxz))} \rightarrow \frac{\text{scope of } \exists y'}{\frac{\text{scope of 2nd } \exists z'}{\exists y(\exists z(Dz \wedge Oyz) \wedge Fyx)}}}}{\exists y(\exists z(Dz \wedge Oxz) \rightarrow \exists z(Dz \wedge Oyz) \wedge Fyx)}}}{\forall x(\exists z(Dz \wedge Oxz) \rightarrow \exists y(\exists z(Dz \wedge Oyz) \wedge Fyx))}$$

Even in this case, however, you might want to choose different variables for every quantifier just as a practical matter, preventing any possibility of confusion.

Sentence **ii** is the trickiest yet. First we paraphrase it as ‘For any x that is a friend of a dog owner, x owns a dog which is also owned by a friend of x ’. Using our stepping-stone tactic, this becomes:

$$\forall x(x \text{ is a friend of a dog owner} \rightarrow x \text{ owns a dog which is owned by a friend of } x).$$

Breaking this down a bit more:

$$\forall x(\exists y(Fxy \wedge y \text{ is a dog owner}) \rightarrow \exists y(Dy \wedge Oxy \wedge y \text{ is owned by a friend of } x)).$$

And a bit more:

$$\forall x(\exists y(Fxy \wedge \exists z(Dz \wedge Oyz)) \rightarrow \exists y(Dy \wedge Oxy \wedge \exists z(Fzx \wedge Ozy))).$$

And we are done!

Practice exercises

A. Using this symbolisation key:

domain: 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

symbolise each of the following sentences in Quantifier:

1. If Cleo loves Bouncer, then Bouncer is a monkey.
2. If both Bouncer and Cleo are alligators, then Amos loves them both.
3. Cleo loves a reptile.
4. Bouncer loves all the monkeys that live at the zoo.
5. All the monkeys that Amos loves love him back.
6. Every monkey that Cleo loves is also loved by Amos.
7. There is a monkey that loves Bouncer, but sadly Bouncer does not reciprocate this love.

B. Using the following symbolisation key:

domain: all animals

Dx : _____ x is a dog

Sx : _____ x likes samurai movies

Lxy : _____ x is larger than _____ y

b : Bertie

e : Emerson

f : Fergis

symbolise the following sentences in Quantifier:

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. No dog is larger than Fergis.
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.

C. Using the following symbolisation key:

domain: 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

symbolise the following English sentences in Quantifier:

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.

D. Using the following symbolisation key:

domain: people

Dx : $\text{---}x$ dances ballet.

Fx : $\text{---}x$ is female.

Mx : $\text{---}x$ is male.

Cxy : $\text{---}x$ is a child of $\text{---}y$.

Sxy : $\text{---}x$ is a sibling of $\text{---}y$.

e : Elmer

j : Jane

p : Patrick

symbolise the following arguments in Quantifier:

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 sons 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 brother who also dances ballet.
12. Every woman who dances ballet is the child of someone who dances ballet.

18 Identity

Consider this sentence:

1. Pavel owes money to everyone.

Let the domain be people; this will allow us to translate ‘everyone’ as a universal quantifier. Offering the symbolisation key:

Oxy : _____ x owes money to _____ y
 p : Pavel

we can symbolise sentence **1** by ‘ $\forall xOpx$ ’. But this has a (perhaps) odd consequence. It requires that Pavel owes money to every member of the domain (whatever the domain may be). The domain certainly includes Pavel. So this entails that Pavel owes money to himself.

Perhaps we meant to say:

2. Pavel owes money to everyone *else*
3. Pavel owes money to everyone *other than* Pavel
4. Pavel owes money to everyone *except* Pavel himself

We want to add something to the symbolisation of **1** to handle these italicised words. Some interesting issues arise as we do so.

18.1 First attempt to handle the argument

The sentences in **2-4** can all be paraphrased as follows:

5. Everyone who isn’t Pavel is such that: Pavel owes money to them.

This is a sentence of the form ‘every F [person who is not Pavel] is G [owed money by Pavel]’. Accordingly it can be symbolised by something with this structure: $\forall xAx \rightarrow Gx$. Here is an attempt to fill in the schematic letters:

Oxy : _____ x owes money to _____ y
 Ixy : _____ x is _____ y
 p : Pavel

With this symbolisation key, here is a proposed symbolisation: $\forall y(Ipy \rightarrow Opy)$.

This symbolisation works well. But, it turns out, it doesn't quite do what we wanted it to. Suppose 'h' names Hikaru, and consider this argument, with the symbolisation next to the English sentences:

6. Pavel owes money to everyone else: $\forall y(\neg Ipy \rightarrow Opy)$.
 7. Hikaru isn't Pavel: $\neg Ihp$.
 So: Pavel owes money to Hikaru: Oph .

This argument is valid in English. But its symbolisation is not valid. If we pick Hikaru to be the value of 'y', we get from 6 the conditional $\neg Iph \rightarrow Oph$. But 7 doesn't give us the antecedent of this conditional: $\neg Iph$ is potentially quite different from $\neg Ihp$. So the argument isn't formally valid.

The argument isn't formally valid, because the sentence ' $\neg Ihp$ ' doesn't formally entail ' $\neg Iph$ ' as a matter of logical structure alone. If the original argument is valid, then we need a symbolisation that as a *matter of logic* allows the distinctness (non-identity) of Hikaru and Pavel to entail the distinctness of Pavel and Hikaru.

18.2 Adding identity

Logicians resolve this issue by adding identity as a new *logical* predicate – one of the structural words with a fixed interpretation.¹ We add a new symbol '=', to clearly differentiate it from our existing predicates.

The symbol '=' is a two-place predicate. Since it is to have a special meaning, we shall write it a bit differently: we put it between two terms, rather than out front. And it *does* have a very particular fixed meaning. Like the quantifiers and connectives, the identity predicate does not need a symbolisation key to fix how it is to be used. Rather, it always gets the same interpretation: $x = y$ always means '_____ x is identical to _____ y '. So identity is a special predicate because it has its meaning as part of logic.

That one thing is identical to another does not mean *merely* that the objects in question are indistinguishable, or that all of the same things are true of them. Rather, it means that the objects in question are *the very same* object.

Now suppose we want to symbolise this sentence:

¹ We don't absolutely have to do this: there are logical languages in which identity is not a logical predicate, and is symbolised by just choosing a two-place predicate like Ixy . But in our logical language Quantifier, we are choosing to treat identity as a structural word.

8. Pavel is Mister Checkov.

Let us add to our symbolisation key:

c : Mister Checkov

Now sentence 8 can be symbolised as ' $p = c$ '. This means that p is c , and it follows that the thing named by ' p ' is the thing named by ' c '.²

We can also now deal with sentences 2–4. All of these sentences can be paraphrased as 'Everyone who isn't Pavel is such that: Pavel owes money to them'. Paraphrasing some more, we get: 'For all x , if x is not Pavel, then x is owed money by Pavel'. Now that we are armed with our new identity symbol, we can symbolise this as ' $\forall x(\neg x = p \rightarrow Opx)$ '.

This last sentence contains the formula ' $\neg x = p$ '. And that might look a bit strange, because the symbol that comes immediately after the ' \neg ' is a variable, rather than a predicate. But this is no problem. We are simply negating the entire formula, ' $x = p$ '. But if this is confusing, you may use this official abbreviation: ' $x \neq y$ ' is the same expression as ' $\neg x = y$ '. I will use both expressions in what follows, but strictly speaking only the latter is officially permitted.

In addition to sentences that use the word 'else', 'other than' and 'except', identity will be helpful when symbolising some sentences that contain the words 'besides' and 'only'. Consider these examples:

9. No one besides Pavel owes money to Hikaru.
10. Only Pavel owes Hikaru money.

Sentence 9 can be paraphrased as, 'No one who is not Pavel owes money to Hikaru'. This can be symbolised by ' $\neg \exists x(\neg x = p \wedge O_x h)$ '. Equally, sentence 9 can be paraphrased as 'for all x , if x owes money to Hikaru, then x is Pavel'. Then it can be symbolised as ' $\forall x(O_x h \rightarrow x = p)$ '.

Sentence 10 can be treated similarly. But there is one subtlety here. Do either sentence 9 or 10 entail that Pavel himself owes money to Hikaru?

18.3 Principles of identity

Return to our argument from premises 6 and 7. We now symbolise it: $\forall y(p \neq y \rightarrow Opy); h \neq p \therefore Oph$. This argument is valid. Given the fixed interpretation we have assigned to ' $=$ ', it is not possible that a first thing be not identical to a second, while the second is identical to the first. So we can conclude, as a matter of logical alone, that $h \neq p$ is equivalent to $p \neq h$.

² One must be careful: the sentence ' $p = c$ ' is, on this symbolisation, about Pavel and Mister Checkov; it is not about 'Pavel' and 'Mister Checkov', which are obviously distinct expressions of English.

This argument rests on one of the logical principles governing identity: that it is SYMMETRIC, so that $x = y$ is equivalent to $y = x$. There are other logical principles of identity too. Everything is itself, and not another thing. So everything has the ‘relation’ of identity to itself: $x = x$ is always true, no matter what name we substitute for x . This is known as REFLEXIVITY of identity. Finally, if x is identical to y , and y identical to z , then it follows that $x = z$. Logicians call this feature of a relation TRANSITIVITY.

The final principle about identity has its own name, LEIBNIZ’ LAW:

If $x = y$ then for any property at all, x has it iff y has it too. That is: every instance of this schematic sentence of Quantifier, for any predicate \mathcal{F} whatsoever, is true:

$$\forall x \forall y (x = y \rightarrow (\mathcal{F}x \leftrightarrow \mathcal{F}y)).$$

Leibniz’ Law certainly entails that identical things are indistinguishable, sharing every property in common. But as we have already noted, identity isn’t merely indistinguishability. Two things might be indistinguishable, but if there are two, they are not strictly identical in the logical sense we are concerned with.

18.4 There are at least ...

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

11. There is at least one apple
12. There are at least two apples
13. There are at least three apples

We shall use the symbolisation key:

Ax : _____ x is an apple

Sentence 11 does not require identity. It can be adequately symbolised by ‘ $\exists xAx$ ’: There is some apple; perhaps many, but at least one.

It might be tempting to also translate sentence 12 without identity. Yet consider the sentence ‘ $\exists x \exists y (Ax \wedge Ay)$ ’. Roughly, this says that there is some apple x in the domain and some apple y in the domain. Since nothing precludes these from being one and the same apple, this would be true even if there were only one apple. In order to make sure that we are dealing with *different* apples, we need an identity predicate. Sentence 12 needs to say that the two apples that exist are not identical, so it can be symbolised by ‘ $\exists x \exists y (Ax \wedge Ay \wedge \neg x = y)$ ’.

Sentence 13 requires talking about three different apples. Now we need three existential quantifiers, and we need to make sure that each will pick out something different: ‘ $\exists x \exists y \exists z (Ax \wedge Ay \wedge Az \wedge x \neq y \wedge y \neq z \wedge x \neq z)$ ’.

18.5 There are at most ...

Now consider these sentences:

14. There is at most one apple.
15. There are at most two apples.

Sentence 14 can be paraphrased as, ‘It is not the case that there are at least *two* apples’. This is just the negation of sentence 12:

$$\neg \exists x \exists y (Ax \wedge Ay \wedge \neg x = y).$$

But sentence 14 can also be approached in another way. It means that if you pick out an object and it’s an apple, and then you pick out an object and it’s also an apple, you must have picked out the same object both times. With this in mind, it can be symbolised by

$$\forall x \forall y ((Ax \wedge Ay) \rightarrow x = y).$$

The two sentences will turn out to be logically equivalent.

In a similar way, sentence 15 can be approached in two equivalent ways. It can be paraphrased as, ‘It is not the case that there are *three* or more distinct apples’, so we can offer:

$$\neg \exists x \exists y \exists z (Ax \wedge Ay \wedge Az \wedge x \neq y \wedge y \neq z \wedge x \neq z).$$

Or, we can read it as saying that if you pick out an apple, and an apple, and an apple, then you will have picked out (at least) one of these objects more than once. Thus:

$$\forall x \forall y \forall z ((Ax \wedge Ay \wedge Az) \rightarrow (x = y \vee x = z \vee y = z)).$$

18.6 There are exactly ...

We can now consider precise statements, like:

16. There is exactly one apple.
17. There are exactly two apples.
18. There are exactly three apples.

Sentence 16 can be paraphrased as, ‘There is *at least* one apple and there is *at most* one apple’. This is just the conjunction of sentence 11 and sentence 14. So we can offer:

$$\exists x Ax \wedge \forall x \forall y ((Ax \wedge Ay) \rightarrow x = y).$$

But it is perhaps more straightforward to paraphrase sentence 16 as, ‘There is a thing x which is an apple, and everything which is an apple is just x itself’. Thought of in this way, we offer:

$$\exists x (Ax \wedge \forall y (Ay \rightarrow x = y)).$$

Similarly, sentence 17 may be paraphrased as, ‘There are *at least* two apples, and there are *at most* two apples’. Thus we could offer

$$\exists x \exists y (Ax \wedge Ay \wedge \neg x = y) \wedge \forall x \forall y \forall z ((Ax \wedge Ay \wedge Az) \rightarrow (x = y \vee x = z \vee y = z)).$$

More efficiently, though, we can paraphrase it as ‘There are at least two different apples, and every apple is one of those two apples’. Then we offer:

$$\exists x \exists y (Ax \wedge Ay \wedge \neg x = y \wedge \forall z (Az \rightarrow (x = z \vee y = z))).$$

Finally, consider these sentence:

19. There are exactly two things.
20. There are exactly two objects.

It might be tempting to add a predicate to our symbolisation key, to symbolise the English predicate ‘_____ is a thing’ or ‘_____ is an object’. But this is unnecessary. Words like ‘thing’ and ‘object’ do not sort wheat from chaff: they apply trivially to everything, which is to say, they apply trivially to every thing. So we can symbolise either sentence with either of the following:

$$\begin{aligned} \exists x \exists y \neg x = y \wedge \neg \exists x \exists y \exists z (\neg x = y \wedge \neg y = z \wedge \neg x = z); \text{ or} \\ \exists x \exists y (\neg x = y \wedge \forall z (x = z \vee y = z)). \end{aligned}$$

Practice exercises

A. Explain why:

- › ‘ $\exists x \forall y (Ay \leftrightarrow x = y)$ ’ is a good symbolisation of ‘there is exactly one apple’.
- › ‘ $\exists x \exists y (\neg x = y \wedge \forall z (Az \leftrightarrow (x = z \vee y = z)))$ ’ is a good symbolisation of ‘there are exactly two apples’.

19 Definite descriptions

In Quantifier, names function rather like names in English. They are simply labels for the things they name, and may be attached arbitrarily, without any indication of the characteristics of what they name.¹

But complex noun phrases can also be used to denote particular things in English, and they do so not merely by acting as arbitrary labels, but often by *describing* the thing they refer to. Consider sentences like:

1. Nick is the traitor.
2. The traitor went to Cambridge.
3. The traitor is the deputy.
4. The traitor is the shortest person who went to Cambridge.

These noun phrases headed by ‘the’ – ‘the traitor’, ‘the deputy’, ‘the shortest person who went to Cambridge’ – are known as DEFINITE DESCRIPTIONS. They are meant to pick out a *unique* object, by using a description which applies to that object and to no other (at least, to no other salient object).

Definite descriptions should be contrasted with *indefinite* descriptions, such as ‘A traitor went to Cambridge’. They should equally be contrasted with *generics*, such as ‘The whale is a mammal’.² The question we face is: can we adequately symbolise definite descriptions in Quantifier?

19.1 Treating definite descriptions as terms

One option would be to introduce new names whenever we come across a definite description. This is not a great idea. We know that *the* traitor – whoever it is – is indeed a

¹ This is not strictly true: consider the name ‘Fido’ which is conventionally the name of a dog. But even here the name doesn’t carry any information in itself about what it names – the fact that we use that as a name only for dogs allows someone who knows that to reasonably infer that Fido is a dog. But ‘Fido is a dog’ isn’t a necessary truth, as it would be if somehow ‘Fido’ carried with it the information that it applies only to dogs.

² It’s inappropriate to ask *which* whale. Moreover the generic use of ‘the whale is a mammal’ can be paraphrased ‘Whales are mammals’. But the fact that ‘The Prime Minister is a Liberal’ cannot be paraphrased as ‘Prime Ministers are Liberals’ indicates the involvement of a genuine definite description.

traitor. We want to preserve that information in our symbolisation. So the symbolisation of ‘The traitor is a traitor’ should be a tautology. But if we symbolise ‘the traitor’ by a name a , the symbolisation will be Ta , which is not a tautology.

A second option would be to use a *new* definite description operator, such as ‘ι’. The idea would be to symbolise ‘the F’ as ‘ ιxFx ’. This is read something like: ‘the thing such that it is F’, though any complex expression with the variable x would work. Expression of the form ιxAx would then behave, grammatically speaking, like names. If we followed this path, then using the following symbolisation key:

domain: people

Tx : _____ x is a traitor

Dx : _____ x is a deputy

Cx : _____ x went to Cambridge

Sxy : _____ x is shorter than _____ y

n : Nick

We could symbolise sentence 1 with ‘ $\iota xTx = n$ ’, sentence 2 with ‘ $C\iota xTx$ ’, sentence 3 with ‘ $\iota xTx = \iota xDx$ ’, and sentence 4 with ‘ $\iota xTx = \iota x(Cx \wedge \forall y((Cy \wedge x \neq y) \rightarrow Sxy))$ ’. This last example may be a bit tricky to parse. In semi-formal English, it says (supposing a domain of persons): ‘the unique person such that they are a traitor is identical with the unique person such that they went to Cambridge and they are shorter than anyone else who went to Cambridge’.

However, even adding this new symbol to our language doesn’t quite help with our initial complaint, since it is not self-evident that the symbolisation of ‘The traitor is a traitor’ as ‘ $T\iota xTx$ ’ yields a tautology. Moreover, it would be nice if we didn’t have to add a new symbol to *Quantifier*. And indeed, we might be able to make do without one.

19.2 Russell’s paraphrase

Bertrand Russell offered an influential account of definite descriptions that might serve our purposes. Very briefly put, he observed that, when we say ‘the F’, where that phrase is a definite description, our aim is to pick out the *one and only* thing that is F (in the appropriate context). Thus Russell offered a systematic paraphrase of sentences involving definite descriptions along these lines:³

the F is G **iff** there is at least one F, *and*
 there is at most one F, *and*
 every F is G

Note a very important feature of this paraphrase: ‘*the*’ *does not appear on the right-side of the equivalence*. This approach would allow us to paraphrase every sentence of the same

³ Bertrand Russell, ‘On Denoting’, 1905, *Mind* 14, pp. 479–93; also Russell, *Introduction to Mathematical Philosophy*, 1919, London: Allen and Unwin, ch. 16.

form as the left hand side into a sentence of the same form as the right hand side, and thus ‘paraphrase away’ the definite description.

It is crucial to notice that we can handle each of the conjuncts of the right hand side of the equivalence in *Quantifier*, using the techniques for dealing with numerical quantification just introduced. We can deal with the three conjuncts on the right-hand side of Russell’s paraphrase as follows:

$$\exists xFx \wedge \forall x\forall y((Fx \wedge Fy) \rightarrow x = y) \wedge \forall x(Fx \rightarrow Gx)$$

In fact, we could express the same point rather more crisply, by recognising that the first two conjuncts just amount to the claim that there is *exactly* one F, and that the last conjunct tells us that that object is F. So, equivalently, we could offer:

$$\exists x(Fx \wedge \forall y(Fy \rightarrow x = y) \wedge Gx)$$

Using these sorts of techniques, we can now symbolise sentences 1–3 without using any new-fangled fancy operator, such as ‘i’.

Sentence 1 is exactly like the examples we have just considered. So we would symbolise it by ‘ $\exists x(Tx \wedge \forall y(Ty \rightarrow x = y) \wedge x = n)$ ’.

Sentence 2 poses no problems either: ‘ $\exists x(Tx \wedge \forall y(Ty \rightarrow x = y) \wedge Cx)$ ’.

Sentence 3 is a little trickier, because it links two definite descriptions. But, deploying Russell’s paraphrase, it can be paraphrased by ‘there is exactly one traitor, x , and there is exactly one deputy, y , and $x = y$ ’. So we can symbolise it by:

$$\exists x\exists y((Tx \wedge \forall z(Tz \rightarrow x = z)) \wedge (Dy \wedge \forall z(Dz \rightarrow y = z)) \wedge x = y)$$

Note that I have made sure that the formula ‘ $x = y$ ’ falls within the scope of both quantifiers!

Let us dispel an initial worry. It seems that I can say ‘the table is brown’ without implying that there is one and only one table in the universe. But doesn’t Russell’s paraphrase literally entail that there is only one table? Indeed it does – it entails that there is only one table *in the domain under discussion*. While sometimes we explicitly restrict the domain, usually we leave it to our background conversational presuppositions to do so (recall §15.5). If I can successfully say ‘the table is brown’ in a conversation with you, for example, some background restriction on the domain must be in place that we both tacitly accept. For example, it might be that the prior discussion has focussed on your dining room, and so the implicit domain is ‘things in that room’. In that case, ‘the table is brown’ is true just in case there is exactly one table in that domain, and it is brown.

19.3 Empty definite descriptions

One of the nice features of Russell’s paraphrase is that it allows us to handle *empty* definite descriptions neatly.

France has no king at present. Now, if we were to introduce a name, 'k', to name the present King of France, then everything would go wrong: remember from §15 that a name must always pick out some object in the domain, and whatever we choose as our domain, it will contain no present kings of France.

Russell's paraphrase neatly avoids this problem. Russell tells us to treat definite descriptions using predicates and quantifiers, instead of names. Since predicates can be empty (see §16), this means that no difficulty now arises when the definite description is empty.

Indeed, Russell's paraphrase helpfully highlights two ways to go wrong in a claim involving a definite description. To adapt an example from Stephen Neale (1990),⁴ suppose I, Antony Eagle, claim:

5. I am grandfather to the present king of France.

Using the following symbolisation key:

a : Antony
 Kx : _____ x is a present king of France
 Gxy : _____ x is a grandfather of _____ y

Sentence 5 would be symbolised by ' $\exists x(\forall y(Ky \leftrightarrow x = y) \wedge Gax)$ '. Now, this sentence isn't true. We might express this by saying: 'Antony isn't the grandfather of the present king of France'. But this sentence is *ambiguous*. The two available readings of this sentence correspond to these two different sentences:

6. There is no one who is both the present King of France and such that Antony is his grandfather.
7. There is a unique present King of France, but Antony is not his grandfather.

Sentence 6 might be paraphrased by 'It is not the case that: Antony is a grandfather of the present King of France'. It will then be symbolised by ' $\neg\exists x(Kx \wedge \forall y(Ky \rightarrow x = y) \wedge Gax)$ '. We might call this *outer* negation, since the negation governs the entire sentence. Note that it will be true if there is no present King of France.

Sentence 7 can be symbolised by ' $\exists x(Kx \wedge \forall y(Ky \rightarrow x = y) \wedge \neg Gax)$ '. We might call this *inner* negation, since the negation occurs within the scope of the definite description. Note that its truth requires that there is a present King of France, albeit one who is a grandchild of Antony.

The idea is: the Russellian paraphrase provides two places for the negation of 'isn't' to slot in: taking scope over the whole sentence, and taking scope just over the predicate ' Gax '. We see evidence for these two options in the ambiguity of the sentence where the role of the negation isn't made clear. One can see the speaker is denying 5, but on what grounds?

⁴ Neale, *Descriptions*, 1990, Cambridge: MIT Press.

Are they accepting that the definite description refers and denying what is said of the present king of France? Or are they denying that it refers?

The operator approach to definite descriptions cannot handle this contrast. There is just one symbolisation of the negation of 5 available in this framework: $\neg \forall x Kx$. The original sentence is false, so this negation must be true. Since sentence 7 is false, this sentence does not express the inner negation of sentence 5. But there is no way to put the negation elsewhere that will express the same claim as 7. ($\forall x \neg Kx$ clearly doesn't do the job – it says Antony is grandfather to the unique thing which isn't the present king of France!) So the sentence 'Antony isn't grandfather to the present king of France' has only one correct symbolisation, and hence only one reading, if the operator approach to definite descriptions is correct. Since it is ambiguous, with multiple readings, the 'λ'-operator approach cannot be correct – it doesn't provide a complex enough grammar for definite description sentences to allow for these kinds of scope ambiguities.

19.4 The adequacy of Russell's paraphrase

The evidence from scope ambiguity we just considered is a substantial point in favour of Russell's paraphrase, at least compared to the operator approach. But how successful is Russell's paraphrase in general? There is a substantial philosophical literature around this issue, but I shall content myself with two observations.

One worry focusses on Russell's treatment of empty definite descriptions. If there are no Fs, then on Russell's paraphrase, both 'the F is G' is and its internal negation, 'the F is non-G', are false. P.F. Strawson suggested that such sentences should not be regarded as false, exactly.⁵ Rather, they involve *presupposition failure*, and he argues that they ought to be regarded as *neither true nor false*, but nevertheless still meaningful.⁶

If we agree with Strawson here, we shall need to revise our logic. For, in our logic, there are only two truth values (True and False), and every meaningful sentence is assigned exactly one of these truth values.⁷

But there is room to disagree with Strawson. Strawson is appealing to some linguistic intuitions, but it is not clear that they are very robust. For example: isn't it just *false*, not 'gappy', that Antony is grandfather to the present King of France? (This is Neale's line.)

⁵ P.F. Strawson, 'On Referring', 1950, *Mind* 59, pp. 320–34.

⁶ A SEMANTIC PRESUPPOSITION of \mathcal{A} is a claim that is a consequence of both \mathcal{A} and $\neg \mathcal{A}$. So 'John has stopped drinking' and its negation 'John hasn't stopped drinking' both entail that John drank, and hence 'John drank' is a semantic presupposition of 'John has stopped drinking'. 'Presupposition failure' occurs when the presupposition of a sentence is false.

⁷ Why? Suppose there is presupposition failure of 'John has stopped drinking', because John never drank. Then 'John has stopped drinking' can't be true since it entails something false, its presupposition 'John drank'. And 'John hasn't stopped drinking' can't be true, since it also entails that same falsehood. So neither can be true. Since one is the negation of the other, if either is false, the other is true. So neither can be false, either. So if there are non-trivial semantic presuppositions – semantic presuppositions that might be false – then we shall have to admit that some meaningful sentences with a false presupposition are neither true nor false. It remains an open question, admittedly, whether presuppositions in the ordinary and intuitive sense are really semantic presuppositions in our sense.

Keith Donnellan raised a second sort of worry, which (very roughly) can be brought out by thinking about a case of mistaken identity.⁸ Two men stand in the corner: a very tall man drinking what looks like a gin martini; and a very short man drinking what looks like a pint of water. Seeing them, Malika says:

8. The gin-drinker is very tall!

Russell's paraphrase will have us render Malika's sentence as:

8'. There is exactly one gin-drinker [in the corner], and whomever is a gin-drinker [in the corner] is very tall.

But now suppose that the very tall man is actually drinking *water* from a martini glass; whereas the very short man is drinking a pint of (neat) gin. By Russell's paraphrase, Malika has said something false. But don't we want to say that Malika has said something *true*?

Again, one might wonder how clear our intuitions are on this case. We can all agree that Malika intended to pick out a particular man, and say something true of him (that he was tall). On Russell's paraphrase, she actually picked out a different man (the short one), and consequently said something false of him. But maybe advocates of Russell's paraphrase only need to explain *why* Malika's intentions were frustrated, and so why she said something false. This is easy enough to do: Malika said something false because she had false beliefs about the men's drinks; if Malika's beliefs about the drinks had been true, then she would have said something true.⁹

To say much more here would lead us into deep philosophical waters. That would be no bad thing, but for now it would distract us from the immediate purpose of learning formal logic. So, for now, we shall stick with Russell's paraphrase of definite descriptions, when it comes to putting things into *Quantifier*. It is certainly the best that we can offer, without significantly revising our logic. And it is quite defensible as an paraphrase.

Practice exercises

A. Using the following symbolisation key:

domain: 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

⁸ Keith Donnellan, 'Reference and Definite Descriptions', 1966, *Philosophical Review* 77, pp. 281-304.

⁹ Interested parties should read Saul Kripke, 'Speaker Reference and Semantic Reference', 1977, in French et al (eds.), *Contemporary Perspectives in the Philosophy of Language*, Minneapolis: University of Minnesota Press, pp. 6-27.

symbolise the following sentences in Quantifier:

1. Hoffhor trusts a vegetarian.
2. Everyone who trusts Ingmar trusts a vegetarian.
3. Everyone who trusts Ingmar trusts someone who trusts a vegetarian.
4. Only Ingmar knows the combination to the safe.
5. Ingmar trusts Hoffhor, but no one else.
6. The person who knows the combination to the safe is a vegetarian.
7. The person who knows the combination to the safe is not a spy.

B. Using the following symbolisation key:

domain: 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.

symbolise each sentence in Quantifier:

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.

C. Using the following symbolisation key:

domain: 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.

symbolise the following sentences in Quantifier:

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; any other creatures in the field must be wingless.
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.

D. In this section, I symbolised 'Nick is the traitor' by ' $\exists x(Tx \wedge \forall y(Ty \rightarrow x = y) \wedge x = n)$ '. Two equally good symbolisations would be:

- › $Tn \wedge \forall y(Ty \rightarrow n = y)$
- › $\forall y(Ty \leftrightarrow y = n)$

Explain why these would be equally good symbolisations.

E. Candace returns to her parents' home to find the family dog with a bandaged nose. Her mother says 'the dog got into a fight with another dog', and this seems a perfectly appropriate thing to say in the circumstances.

Does this example pose a problem for Russell's approach to definite descriptions?

F. Some people have argued that the following two sentences are ambiguous. Are they? If they are, explain how this fact might be used to provide support to Russell's paraphrases of definite description sentences.

1. The prime minister has always been Australian.
2. The number of planets is necessarily eight.

G. Russell's paraphrase of an indefinite description sentence like 'I met a man' is: *there is at least one thing x such that x is male and human and I met x .*

Note that the word 'is' has apparently two readings: sometimes, as in 'Fido is heavy', it indicates *predication*; sometimes, as in 'Fido is Rover', it indicates identity (in the case, the same dog is known by two names). Something interesting arises if Russell's account of indefinite descriptions is right, since in an example like 'Fido is a dog of unusual size', we might interpret the 'is' in either way, roughly:

1. Fido is identical to a dog of unusual size;
2. Fido has the property of being a dog of unusual size.

Suppose that ' Ux ' symbolises the property of being a dog of unusual size, then our two readings can be symbolised ' $\exists x(Ux \wedge f = x)$ ' and ' $U(f)$ '.

Is there any significant difference in meaning between these two symbolisations?

20 Sentences of Quantifier

We know how to represent English sentences in Quantifier. The time has finally come to define the notion of a *sentence* of Quantifier.

20.1 Expressions

There are six kinds of symbols in Quantifier:

Predicates	A, B, C, \dots, Z
with subscripts, as needed	$A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \dots$
Names	a, b, c, \dots, r
with subscripts, as needed	$a_1, b_{224}, h_7, m_{32}, \dots$
Variables	s, t, u, v, w, x, y, z
with subscripts, as needed	$x_1, y_1, z_1, x_2, \dots$
Connectives	$\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
Brackets	$(,)$
Quantifiers	\forall, \exists

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

20.2 Terms and formulas

In §6, we went straight from the statement of the vocabulary of Sentential to the definition of a sentence of Sentential. In Quantifier, we shall have to go via an intermediary stage: via the notion of a *formula*. The intuitive idea is that a formula is any sentence, or

anything which can be turned into a sentence by adding quantifiers out front. But this will take some unpacking.

We start by defining the notion of a term.

A TERM is any name or any variable.

So, here are some terms:

$$a, b, x, x_1x_2, y, y_{254}, z$$

We next need to define atomic formulas.

1. If \mathcal{R} is any n -place predicate other than the identity predicate '=', and t_1, t_2, \dots, t_n are terms (not necessarily distinct from one another), then $\mathcal{R}t_1t_2\dots t_n$ is an atomic formula.
2. If t_1 and t_2 are terms, then $t_1 = t_2$ is an atomic formula.
3. Nothing else is an atomic formula.

The use of script fonts here follows the conventions laid down in §7. So, ' \mathcal{R} ' is not itself a predicate of Quantifier. Rather, it is a symbol of our metalanguage (augmented English) that we use to talk about any predicate of Quantifier. Similarly, ' t_1 ' is not a term of Quantifier, but a symbol of the metalanguage that we can use to talk about any term of Quantifier. So, where ' F ' is a one-place predicate, ' G ' is a three-place predicate, and ' S ' is a six-place predicate, here are some atomic formulas:

$$\begin{aligned} x &= a \\ a &= b \\ Fx \\ Fa \\ Gxay \\ Gaaa \\ Sx_1x_2abyx_1 \\ Sby_{254}zaaz \end{aligned}$$

Remember that we allow 0-place predicates too; these predicates, when followed by zero terms, are atomic formulae of Quantifier.

Once we know what atomic formulas are, we can offer recursion clauses to define arbitrary formulas. The first few clauses are exactly the same as for Sentential.

1. Every atomic formula is a formula.
2. If \mathcal{A} is a formula, then $\neg\mathcal{A}$ is a formula.
3. If \mathcal{A} and \mathcal{B} are formulas, then $(\mathcal{A} \wedge \mathcal{B})$ is a formula.
4. If \mathcal{A} and \mathcal{B} are formulas, then $(\mathcal{A} \vee \mathcal{B})$ is a formula.
5. If \mathcal{A} and \mathcal{B} are formulas, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a formula.
6. If \mathcal{A} and \mathcal{B} are formulas, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a formula.
7. If \mathcal{A} is a formula and x is a variable, then $\forall x\mathcal{A}$ is a formula.
8. If \mathcal{A} is a formula and x is a variable, then $\exists x\mathcal{A}$ is a formula.
9. Nothing else is a formula.

So, assuming again that ‘ F ’ is a one-place predicate, ‘ G ’ is a three-place predicate and ‘ H ’ is a six place-predicate, here are some formulas:

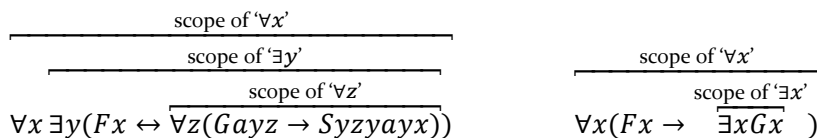
$$\begin{aligned}
 &Fx \\
 &Gayz \\
 &Syzyayx \\
 &(Gayz \rightarrow Syzyayx) \\
 &\forall z(Gayz \rightarrow Syzyayx) \\
 &Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx) \\
 &\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx)) \\
 &\forall x\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx)) \\
 &\forall x(Fx \rightarrow \exists xGx)
 \end{aligned}$$

We can now give a formal definition of scope, which incorporates the definition of the scope of a quantifier. Here we follow the case of Sentential, though we note that a connective can be either a connective or a quantifier:

The MAIN CONNECTIVE in a formula is the operator that was introduced last, when that formula was constructed using the recursion rules.

The SCOPE of a connective in a formula is the subformula for which that operator is the main connective.

So we can graphically illustrate the scope of the quantifiers in the last two examples thus:



Note that it follows from our recursive definition that ‘ $\forall xFa$ ’ is a formula. While puzzling on its face, with that quantifier governing no variable in its scope, it nevertheless is a formula of the language. It turns out that, with that redundant quantifier, it is equivalent to ‘ Fa ’. Eliminating such formulae from the language involves greatly increased complication in the definition of a formula for no important gain.

20.3 Sentences

Recall that we are largely concerned in logic with assertoric sentences: sentences that can be either true or false. Many formulas are not sentences. Consider the following symbolisation key:

domain: people
 Lxy : _____ x loves _____ y
 b : Boris

Consider the atomic formula ‘ Lzz ’. All atomic formula are formulas, so ‘ Lzz ’ is a formula. But can it be true or false? You might think that it will be true just in case the person named by ‘ z ’ loves herself, in the same way that ‘ Lbb ’ is true just in case Boris (the person named by ‘ b ’) loves himself. *But ‘ z ’ is a variable, and does not name anyone or any thing.* It is true that we can sometimes manage to make a claim by saying ‘it loves it’, the best English rendering of ‘ Lzz ’. But we can only do so by making use of contextual cues to supply referents for the pronouns ‘it’ – contextual cues that artificial languages like Quantifier lack.

Of course, if we put an existential quantifier out front, obtaining ‘ $\exists zLzz$ ’, then this would be true iff someone loves themselves. Equally, if we wrote ‘ $\forall zLzz$ ’, this would be true iff everyone loves themselves. The point is that we need a quantifier to tell us how to deal with a variable.

Let’s make this idea precise.

A BOUND VARIABLE is an occurrence of a variable x that is within the scope of either $\forall x$ or $\exists x$.
 A FREE VARIABLE is any variable that is not bound.

For example, consider the formula

$$\forall x(Ex \vee Dy) \rightarrow \exists z(Ez \rightarrow Lzx)$$

The scope of the universal quantifier ‘ $\forall x$ ’ is ‘ $\forall x(Ex \vee Dy)$ ’, so the first ‘ x ’ is bound by the universal quantifier. However, the second and third occurrence of ‘ x ’ are free. Equally, the ‘ y ’ is free. The scope of the existential quantifier ‘ $\exists z$ ’ is ‘ $(Ez \rightarrow Lzx)$ ’, so ‘ z ’ is bound.

In our last example from the previous section, ‘ $\forall x(Fx \rightarrow \exists xGx)$ ’, the variable ‘ x ’ in ‘ Gx ’ is bound by the quantifier ‘ $\exists x$ ’, and so x is free in ‘ $Fx \rightarrow \exists xGx$ ’ only when it appears in

' Fx '. So while the scope of ' $\forall x$ ' is the whole sentence, it nevertheless doesn't bind every variable in its scope – only those such that, were it absent, would be free. (So we might say an occurrence of a variable x is BOUND BY an occurrence of quantifier $\forall x/\exists x$ just in case it would have been free had that quantifier been omitted.)

Finally we can say the following.

The SENTENCES of Quantifier are just exactly those formulae of Quantifier that contain no free variables.

Since an atomic formula formed by a 0-place predicate contains no terms at all, and hence cannot contain a variable, every such expression is a sentence – they are just the atomic sentences of Sentential. Any other formula which contains no variables, but only names, is also a sentence, as well as all those formulae which contain only bound variables.

20.4 Bracketing conventions

We will adopt the same notational conventions governing brackets that we did for Sentential (see §6 and §10.3.)

First, we may omit the outermost brackets of a formula, and we sometimes use variably sized brackets to aid with readability.

Second, we may omit brackets between each pair of conjuncts when writing long series of conjunctions.

Third, we may omit brackets between each pair of disjuncts when writing long series of disjunctions.

Practice exercises

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

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

Chapter 5

Interpretations

21 Extensionality

Recall that `Sentential` is a truth-functional language. Its connectives are all truth-functional, and *all* that we can do with `Sentential` is key sentences to particular truth values. We can do this *directly*. For example, we might stipulate that the `Sentential` sentence ' P ' is to be true. Alternatively, we can do this *indirectly*, offering a symbolisation key, e.g.:

P : Big Ben is in London

But recall from §9 that this should be taken to mean:

- › The `Sentential` sentence ' P ' is to take the same truth value as the English sentence 'Big Ben is in London' (whatever that truth value may be)

The point that I emphasised is that `Sentential` cannot handle differences in meaning that go beyond mere differences in truth value.

21.1 Symbolising versus translating

`Quantifier` has some similar limitations. It gets beyond mere truth values, since it enables us to split up sentences into terms, predicates and quantifier expressions. This enables us to consider what is *true of* some particular object, or of some or all objects. But we can do no more than that.

When we provide a symbolisation key for some `Quantifier` predicates, such as:

Cx : _____ x lectures on logic in Adelaide in Semester 2, 2017

we do not carry the *meaning* of the English predicate across into our `Quantifier` predicate. We are simply stipulating something like the following:

- › ' Cx ' and '_____ x lectures on logic in Adelaide in Semester 2, 2017' are to be *true of* exactly the same things.

So, in particular:

- › ‘ Cx ’ is to be true of all and only those things which lecture on logic in Adelaide in Semester 2, 2017 (whatever those things might be).

This is an indirect stipulation. Alternatively we can stipulate predicate extensions directly. We can stipulate that ‘ Cx ’ is to be true of Antony Eagle and Jon Opie, and only them. As it happens, this direct stipulation would have the same effect as the indirect stipulation. But note that the English predicates ‘_____ is Antony Eagle or Jon Opie’ and ‘_____ lectures on logic in Adelaide in Semester 2, 2017’ have very different meanings!

The point is that Quantifier does not give us any resources for dealing with nuances of meaning. When we interpret Quantifier, all we are considering is what the predicates are actually true of. For this reason, I say only that Quantifier sentences *symbolise* English sentences. It is doubtful that we are *translating* English into Quantifier, for translations should preserve meanings.

This is normally summed up by saying that Quantifier is an EXTENSIONAL LANGUAGE. The extension of an English expression is just the things to which it actually applies. So the extension of a name is the thing actually named; and the extension of a predicate is just those things it actually covers. An extensional language is one where expressions with the same extension are *synonyms*. English is not an extensional language. While the name ‘Malcolm Turnbull’ and the definite noun phrase ‘the Prime Minister’ have the same extension, they are not synonyms. For one can replace an expression in a sentence by one of its synonyms, and the resulting sentence will mean the same thing as the original - and will have the same truth value. But while ‘Malcolm Turnbull has always been Malcolm Turnbull’ is true, ‘Malcolm Turnbull has always been the Prime Minister’ is not true. Despite having the same extension, these expressions differ in meaning. In Quantifier, by contrast, the substitution of one name for another with the same extension will always yield a sentence with the same truth-value; likewise with the substitution of one predicate for another with the same extension.¹

¹ What this shows, in passing, is that Quantifier lacks the resources to express things like ‘_____ has always been _____’ or ‘_____ is necessarily _____’, which would allow us to separate non-synonymous expressions with the same present extension.

Another construction with apparently similar effects is when a true identity, such as ‘Lewis Carroll is Charles Lutwidge Dodgson’, is embedded in a belief report such as ‘AE believes that Lewis Carroll is Charles Lutwidge Dodgson’. The belief report appears to be false, if AE doesn’t know that ‘Lewis Carroll’ is a pen name for the Oxford mathematician. But still, this seems hard to deny: ‘AE believes that Lewis Carroll is Lewis Carroll’. Many have used this to argue that the meaning of a name in English is not just its extension. But this is actually a rather controversial case, unlike the example in the main text: many philosophers think that the meaning of a name in English just is its extension. But let us be clear: no one generalises this to predicates. Everyone agrees that the meaning of an English predicate is not just its extension.

21.2 A word on extensions

We can stipulate directly what predicates are to be true of. So it is worth noting that our stipulations can be as arbitrary as we like. For example, we could stipulate that ‘ Hx ’ should be true of, and only of, the following objects:

David Cameron
the number π
every top-F key on every piano ever made

Now, the objects that we have listed have nothing particularly in common. But this doesn’t matter. Logic doesn’t care about what strikes us mere humans as ‘natural’ or ‘similar’. Armed with this interpretation of ‘ Hx ’, suppose I now add to my symbolisation key:

d : David Cameron
 n : Nick Clegg
 p : the number π

Then ‘ Hd ’ and ‘ Hp ’ will both be true, on this interpretation, but ‘ Hn ’ will be false, since Nick Clegg was not among the stipulated objects.

This process of explicit stipulation is just giving the extension of a predicate by a list of items falling under it. The much more common way of identifying the extension of a predicate is to derive it from a classification rule that sorts everything into those items that fall under it and those that do not. Such a classification rule is often what people think of as the meaning of a predicate.

21.3 Many-place predicates

All of this is quite easy to understand when it comes to one-place predicates. But it gets much messier when we consider two-place predicates. Consider a symbolisation key like:

Lxy : _____ x loves _____ y

Given what I said above, this symbolisation key should be read as saying:

- ‘ Lxy ’ and ‘_____ x loves _____ y ’ are to be true of exactly the same things

So, in particular:

- ‘ Lxy ’ is to be true of x and y (in that order) iff x loves y .

It is important that we insist upon the order here, since love – famously – is not always reciprocated. (Note that ‘ x ’ and ‘ y ’ here are symbols of augmented English, and that they

are being *used*. By contrast, ‘*x*’ and ‘*y*’ are symbols of Quantifier, and they are being *mentioned*.)

That is an indirect stipulation. What about a direct stipulation? This is slightly harder. If we *simply* list objects that fall under ‘*Lxy*’, we will not know whether they are the lover or the beloved (or both). We have to find a way to include the order in our explicit stipulation.

To do this, we can specify that two-place predicates are true of *pairs* of objects, where the order of the pair is important. Thus we might stipulate that ‘*Bxy*’ is to be true of, and only of, the following pairs of objects:

⟨Lenin, Marx⟩
 ⟨Heidegger, Sartre⟩
 ⟨Sartre, Heidegger⟩

Here the angle-brackets keep us informed concerning order. Suppose I now add the following stipulations:

l: Lenin
m: Marx
h: Heidegger
s: Sartre

Then ‘*Blm*’ will be true, since ⟨Lenin, Marx⟩ was in my explicit list. But ‘*Bml*’ will be false, since ⟨Marx, Lenin⟩ was not in my list. However, both ‘*Bhs*’ and ‘*Bsh*’ will be true, since both ⟨Heidegger, Sartre⟩ and ⟨Sartre, Heidegger⟩ are in my explicit list

To make these ideas more precise, we would need to develop some *set theory*, the mathematical theory of collections. This would give you some tools for modelling extensions and ordered pairs (and ordered triples, etc.). Indeed, set theoretic tools lie at the heart of contemporary linguistic approaches to meaning in natural languages. However, we have neither the time nor the need to cover set theory in this book. I shall leave these notions at an imprecise level; the general idea will be clear enough, I hope.

21.4 Semantics for identity

Identity is a special predicate of Quantifier. We write it a bit differently than other two-place predicates: ‘*x = y*’ instead of ‘*Ixy*’ (for example). More important, though, its meaning is fixed as to be genuine *identity*, once and for all. Given a domain, the extension of ‘=’ comprises just those pairs consisting of any member of the domain and itself. If the domain is numbers, for example, the extension of ‘=’ on this domain will be

⟨1,1⟩, ⟨2,2⟩, ⟨3,3⟩,

If two names *a* and *b* are assigned to the same object in a given symbolisation key, and thus have the same extension, ‘*a = b*’ will also be true on that symbolisation key. And

since the two names have the same extension, and *Quantifier* is extensional, substituting one name for another will not change the truth value of any *Quantifier* sentence. So, in particular, if ‘*a*’ and ‘*b*’ name the same object, then all of the following will be true:

$$\begin{aligned} Aa &\leftrightarrow Ab \\ Ba &\leftrightarrow Bb \\ Raa &\leftrightarrow Rbb \\ Raa &\leftrightarrow Rab \\ Rca &\leftrightarrow Rcb \\ \forall xRxa &\leftrightarrow \forall xRxb \end{aligned}$$

This fact is sometimes called the *INDISCERNIBILITY OF IDENTICALS* – if two names denote the same thing, then no claim will be true when formulated using one of those names but false when formulated using the other in its place. That is because, at the level of extensions, there really is just one thing that those claims are about, and how could that one thing be discernible from itself?

The converse claim – the *IDENTITY OF INDISCERNIBLES* – is much more controversial. It is not defensible in *Quantifier*. Suppose two objects differ in some property *P*, but our symbolisation key includes no predicate which denotes *P*. Then we might not have any sentence true of the one object and false of the other, even though they are *distinct* objects which have different properties. We can’t tell them apart, not because they are really just one, but because we don’t have the right words to talk about how they are different. *Quantifier* thus allows that exactly the same predicates might be true of two distinct objects, and thus they are distinct but indiscernible when restricted to the resources of our symbolisation.

21.5 Interpretations

I defined a *valuation* in *Sentential* as any assignment of truth and falsity to atomic sentences. In *Quantifier*, I am going to define an *INTERPRETATION* as consisting of three things:

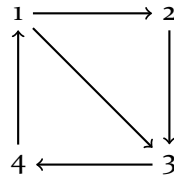
- › the specification of a domain;
- › for each name that we care to consider, an assignment of exactly one object within the domain as its extension;
- › for each non-logical predicate (i.e., any predicate other than ‘=’) that we care to consider, a specification of its extension: what things (or pairs of things, or triples of thing, etc.) the predicate is to be true of.

A 0-place predicate (what we called ‘atomic sentences’ in *Sentential*) has a truth value, either T or F, as its extension – thus the valuations of *Sentential* are interpretations of

a very restricted part of Quantifier— those interpretations which only care to consider 0-place predicates.

The symbolisation keys that I considered in chapter 4 consequently give us one very convenient way to present an interpretation. We shall continue to use them throughout this chapter. It is important to note, however, that the symbolisation key tends to associate each predicate with a *property*, and then uses that property as a classification rule to determine the extension of the predicate. We can, and sometimes do, offer an interpretation by simply giving the extensions of the interpreted Quantifier expressions directly.

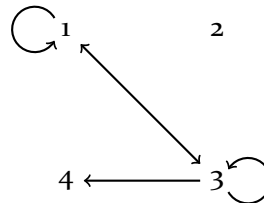
One way of presenting an interpretation that is often convenient is to present its extensions *diagrammatically*. Suppose we want to consider just a single two-place predicate, ' Rxy '. Then we can represent it just by drawing an arrow between two objects, and stipulate that ' Rxy ' is to hold of x and y just in case there is an arrow running from x to y in our diagram. As an example, we might offer:



This would be suitable to characterise an interpretation whose domain is the first four positive whole numbers, and which interprets ' Rxy ' as being true of and only of:

$\langle 1, 2 \rangle, \langle 2, 3 \rangle, \langle 3, 4 \rangle, \langle 4, 1 \rangle, \langle 1, 3 \rangle$

Equally we might offer:



for an interpretation with the same domain, which interprets ' Rxy ' as being true of and only of:

$\langle 1, 3 \rangle, \langle 3, 1 \rangle, \langle 3, 4 \rangle, \langle 1, 1 \rangle, \langle 3, 3 \rangle$

If we wanted, we could make our diagrams more complex. For example, we could add names as labels for particular objects. Equally, to symbolise the extension of a one-place predicate, we might simply draw a ring around some particular objects and stipulate that the thus encircled objects (and only them) are to fall under the predicate ' Hx ', say.

22 Truth in Quantifier

We know what interpretations are. Since, among other things, they tell us which predicates are true of which objects, they will provide us with an account of the truth of atomic sentences. But we must also present a detailed account of what it is for an arbitrary Quantifier sentence to be true or false in an interpretation.

We know from §20 that there are three kinds of sentence in Quantifier:

- › atomic sentences (i.e., atomic formulae of Quantifier which have no free variables);
- › sentences whose main connective is a sentential connective; and
- › sentences whose main connective is a quantifier.

We need to explain truth for all three kinds of sentence.

I shall offer a completely general explanation in this section. However, to try to keep the explanation comprehensible, I shall at several points use the following interpretation:

domain: all people born before 2000CE

a : Aristotle

b : Bush

Wx : _____ x is wise

Rxy : _____ x was born before _____ y

This will be my *go-to example* in what follows.

22.1 Atomic sentences

The truth of atomic sentences should be fairly straightforward. The sentence ' Wa ' should be true just in case ' Wx ' is true of ' a '. Given our go-to interpretation, this is true iff Aristotle is wise. Aristotle is wise. So the sentence is true. Equally, ' Wb ' is false on our go-to interpretation.

Likewise, on this interpretation, ' Rab ' is true iff the object named by ' a ' was born before the object named by ' b '. Well, Aristotle was born before Bush. So ' Rab ' is true. Equally, ' Raa ' is false: Aristotle was not born before Aristotle.

Dealing with atomic sentences, then, is very intuitive. When \mathcal{R} is an n -place predicate and a_1, a_2, \dots, a_n are names,

$\mathcal{R}a_1a_2\dots a_n$ is true in an interpretation **iff**
 \mathcal{R} is true of the objects named by a_1, a_2, \dots, a_n in that interpretation (considered in that order).
 A 0-place predicate is ‘true of’ its accompanying zero objects iff it is true.

Here in fact, we are using Sentential valuations as constituents of Quantifier interpretations: a valuation just fixes the interpretation of the relevant 0-place predicates in an interpretation.

Recall that there is a second kind of atomic sentence, recall: two names connected by an identity sign constitute an atomic sentence. This kind of atomic sentence is also easy to handle. Where a and b are any names,

$a = b$ is true in an interpretation **iff**
 a and b name the very same object in that interpretation

So in our go-to interpretation, ‘ $a = b$ ’ is false, since Aristotle is distinct from Bush.

22.2 Sentential connectives

We saw in §20 that Quantifier sentences can be built up from simpler ones using the truth-functional connectives that were familiar from Sentential. The rules governing these truth-functional connectives are *exactly* the same as they were when we considered Sentential. Here they are:

Where \mathcal{A} and \mathcal{B} are any sentences of Quantifier,

- › $\mathcal{A} \wedge \mathcal{B}$ is true in an interpretation **iff**
 both \mathcal{A} is true and \mathcal{B} is true in that interpretation
- › $\mathcal{A} \vee \mathcal{B}$ is true in an interpretation **iff**
 either \mathcal{A} is true or \mathcal{B} is true in that interpretation
- › $\neg\mathcal{A}$ is true in an interpretation **iff**
 \mathcal{A} is false in that interpretation
- › $\mathcal{A} \rightarrow \mathcal{B}$ is true in an interpretation **iff**
 either \mathcal{A} is false or \mathcal{B} is true in that interpretation
- › $\mathcal{A} \leftrightarrow \mathcal{B}$ is true in an interpretation **iff**
 \mathcal{A} has the same truth value as \mathcal{B} in that interpretation

This presents the very same information as the schematic truth tables for the connectives; it just does it in a slightly different way. Some examples will probably help to illustrate the idea. On our go-to interpretation:

- ' $a = a \wedge Wa$ ' is true
- ' $Rab \wedge Wb$ ' is false because, although ' Rab ' is true, ' Wb ' is false
- ' $a = b \vee Wa$ ' is true
- ' $\neg a = b$ ' is true
- ' $Wa \wedge \neg(a = b \wedge Rab)$ ' is true, because ' Wa ' is true and ' $a = b$ ' is false

Make sure you understand these examples.

22.3 When the main connective is a quantifier

The exciting innovation in *Quantifier*, though, is the use of *quantifiers*. And in fact, expressing the truth conditions for quantified sentences is a bit more fiddly than one might expect.

Here is a naïve first thought. We want to say that ' $\forall xFx$ ' is true iff ' Fx ' is true of everything in the domain. This should not be too problematic: our interpretation will specify directly what ' Fx ' is true of.

Unfortunately, this naïve first thought is not general enough. For example, we want to be able to say that ' $\forall x\exists yLxy$ ' is true just in case ' $\exists yLxy$ ' is true of everything in the domain. And this is problematic, since our interpretation does not directly specify what ' $\exists yLxy$ ' is to be true of. Instead, whether or not this is true of something should follow just from the interpretation of ' Lxy ', the domain, and the meanings of the quantifiers.

So here is a naïve second thought. We might try to say that ' $\forall x\exists yLxy$ ' is to be true in an interpretation iff $\exists yLay$ is true for *every* name a that we have included in our interpretation. And similarly, we might try to say that $\exists yLay$ is true just in case Lab is true for *some* name b that we have included in our interpretation.

Unfortunately, this is not right either. To see this, observe that in our go-to interpretation, we have only given interpretations for *two* names, ' a ' and ' b '. But the domain – all people born before the year 2000CE – contains many more than two people. I have no intention of trying to name *all* of them!

So here is a third thought. (And this thought is not naïve, but correct.) Although it is not the case that we have named *everyone*, each person *could* have been given a name. So we should focus on this possibility of extending an interpretation, by adding a new name. I shall offer a few examples of how this might work, centring on our go-to interpretation, and I shall then present the formal definition.

In our go-to interpretation, ' $\exists xRbx$ ' should be true. After all, in the domain, there is certainly someone who was born after Bush. Lady Gaga is one of those people. Indeed, if we were to extend our go-to interpretation – temporarily, mind – by adding the name ' c ' to

refer to Lady Gaga, then $'Rbc'$ would be true on this extended interpretation. And this, surely, should suffice to make $'\exists xRbx'$ true on the original go-to interpretation.

In our go-to interpretation, $'\exists x(Wx \wedge Rxa)'$ should also be true. After all, in the domain, there is certainly someone who was both wise and born before Aristotle. Socrates is one such person. Indeed, if we were to extend our go-to interpretation by letting a new name, $'c'$, denote Socrates, then $'Wc \wedge Rca'$ would be true on this extended interpretation. Again, this should surely suffice to make $'\exists x(Wx \wedge Rxa)'$ true on the original go-to interpretation.

In our go-to interpretation, $'\forall x\exists yRxy'$ should be false. After all, consider the last person born in the year 1999. I don't know who that was, but if we were to extend our go-to interpretation by letting a new name, $'d'$, denote that person, then we would not be able to find anyone else in the domain to denote with some further new name, perhaps $'e'$, in such a way that $'Rde'$ would be true. Indeed, no matter *whom* we named with $'e'$, $'Rde'$ would be false. And this observation is surely sufficient to make $'\exists yRdy'$ false in our extended interpretation. And this is sufficient to make $'\forall x\exists yRxy'$ false on the original interpretation.

If you have understood these three examples, then that's what matters. Strictly speaking, though, we still need to give a precise definition of the truth conditions for quantified sentences. The result, sadly, is a bit ugly, and requires a few new definitions. Brace yourself!

Suppose that \mathcal{A} is a formula, x is a variable, and c is a name. We shall write $'\mathcal{A}|_{c \sim x}'$ to represent the formula that results from replacing every free occurrence of x in \mathcal{A} by c . So if we began with the formula $'Fxy'$, the metalanguage expression $'Fxy'|_{c \sim x}$ denotes the formula $'Fcy'$. If we began with $'Fyy'$, then $'Fyy'|_{c \sim x}$ just is the original formula $'Fyy'$, since there are no instances of $'x'$ in $'Fyy'$ to be replaced. If we began with $'Fx \vee \forall xGx'$, then $'Fx \vee \forall xGx'|_{c \sim x}$ is $'Fc \vee \forall xGx'$, since neither occurrence of $'x'$ in $'\forall xGx'$ is free.

Suppose we begin with a quantified formula, $\forall x\mathcal{A}$ or $\exists x\mathcal{A}$. If we strip off the quantifier, and pick any name c , then $\mathcal{A}|_{c \sim x}$ is known as a SUBSTITUTION INSTANCE of the original quantified formulae, and c may be called the INSTANTIATING NAME. So:

$$\exists x(Rex \leftrightarrow Fx)$$

is a substitution instance of

$$\forall y\exists x(Ryx \leftrightarrow Fx)$$

with the instantiating name $'e'$, because $'\exists x(Ryx \leftrightarrow Fx)'|_{e \sim y}$ turns out to be $'\exists x(Rex \leftrightarrow Fx)'$.

Armed with this notation, the rough idea is as follows. The sentence $\forall x\mathcal{A}$ will be true iff $\mathcal{A}|_{c \sim x}$ is true no matter what object (in the domain) we name with c . Similarly, the sentence $\exists x\mathcal{A}$ will be true iff there is *some* way to assign the name c to an object that makes $\mathcal{A}|_{c \sim x}$ true. More precisely, we stipulate:

- › $\forall x\mathcal{A}$ is true in an interpretation **iff**
 $\mathcal{A}|_{c\rightsquigarrow x}$ is true in *every* interpretation that extends the original interpretation by assigning an object to some name c not appearing in \mathcal{A} (without changing the original interpretation in any other way).
- › $\exists x\mathcal{A}$ is true in an interpretation **iff**
 $\mathcal{A}|_{c\rightsquigarrow x}$ is true in *some* interpretation that extends the original interpretation by assigning an object to some name c not appearing in \mathcal{A} (without changing the original interpretation in any other way).

That is: we pick a new name that doesn't appear in ' \mathcal{A} '.¹ We uniformly replace any occurrences of the variable ' x ' in ' \mathcal{A} ' by our new name, which creates a substitution instance of ' $\forall x\mathcal{A}$ ' and ' $\exists x\mathcal{A}$ '. Then if this substitution instance is true on every (respectively, some) way of adding an interpretation of the new name to our existing interpretation, then ' $\forall x\mathcal{A}$ ' (respectively, ' $\exists x\mathcal{A}$ ') is true on that existing interpretation.

To be clear: all this is doing is formalising (very pedantically) the intuitive idea expressed above. The result is a bit ugly, and the final definition might look a bit opaque. Hopefully, though, the *spirit* of the idea is clear.

22.4 Satisfaction

The discussion of truth in §§22.1–22.3 only ever involved assigning truth values to sentences. When confronted with formulae involving free variables, we altered them by substituting a new name for the variable, converting it into a sentence, and temporarily extending the interpretation to cover that new name. This is a significant departure from the definition of truth for Sentential sentences in §8. There, we showed how the truth value of a complex Sentential sentence in a valuation depended on the truth value of its constituents in that same valuation. By contrast, on the approach just outlined, the truth value of ' $\exists xFx$ ' in an interpretation depends on the truth value of some other sentence ' Fc ', which is not a constituent of ' $\exists xFx$ ', in some different (although related) interpretation!

There is another way to proceed, which allows arbitrary formulae of Quantifier, even those with free variables, to be assigned (temporarily) a truth value. This approach can let the truth value of ' $\exists xFx$ ' in an interpretation depend on the temporary truth value of ' Fx ' in that same interpretation. This is conceptually neater than the approach just introduced, and I present it briefly here as an alternative to the substitutional approach of the preceding sections. *This section should be regarded as optional, and only for the dedicated student of logic.*

The inspiration for the approach comes, once again, from thinking of variables in Quantifier as behaving like pronouns in English. In §15.4 we gave a gloss of 'Someone

¹ There will always be such a new name: any given sentence of Quantifier only contains finitely many names, but Quantifier has a potentially infinite stock of names to draw from.

is angry' as 'Some person is such that: they are angry'. Concentrate on 'they are angry'. This sentence featuring a bare pronoun doesn't express any specific claim intrinsically. But we can, temporarily, fix a referent for the pronoun 'they' – temporarily elevating it to the status of a name. If we do so, the sentence can be evaluated. We can introduce a referent by pointing: 'They [points to someone] are angry'. Or we can fix a referent by simply assigning one: 'Consider that person over there. They are angry'. If we can find someone or other to temporarily be the referent of the pronoun 'they', then it will be true that there is someone such that they are angry. If no matter who we fix as the referent, they are angry, then it will be true that everyone is such that they are angry.

This is, in a nutshell, the idea we will use to handle quantification in *Quantifier*. Let us introduce some terminology:

A **VARIABLE ASSIGNMENT** over an interpretation is an assignment of exactly one object from the domain of that interpretation to each variable that we care to consider.

If we have an interpretation, and a variable assignment, then we can evaluate every formula of *Quantifier* – not only sentences. Of course, the evaluation of the open formulae will be very fragile, since even given a single background interpretation, a formula like ' Fx ' might be true relative to one variable assignment and false relative to another.

Let us start, as before, by giving rules for evaluating atomic formulae of *Quantifier*, given an interpretation and a variable assignment.

Where \mathcal{R} is any n -place predicate ($n \geq 0$), and t_1, \dots, t_n are any terms – variables or names, then:

- › $\mathcal{R}t_1 \dots t_n$ is true on a variable assignment over an interpretation **iff** \mathcal{R} is true of the objects assigned to the terms t_1, \dots, t_n by that variable assignment and that interpretation, considered in that order.
- › $t_1 = t_2$ is true on a variable assignment over an interpretation **iff** t_1 and t_2 are assigned the very same object by that variable assignment in that interpretation.

These are very similar clauses to those we saw for atomic sentences in §22.3. Indeed, when we are considering atomic sentences of *Quantifier*, the mention of a variable assignment is redundant, since no atomic sentence contains a variable. (If an atomic formula contains a variable, the variable would be free and the formula thus not a sentence.)

The recursion clauses that extend truth for atomic formulae to truth for arbitrary formulae are these (I omit the clauses for \forall , \rightarrow , and \leftrightarrow , which you can easily fill in yourself, following the model in §22.2):

Where \mathcal{A} and \mathcal{B} are formulae of Quantifier, and x is a variable:

- › $\neg\mathcal{A}$ is true on a variable assignment over an interpretation **iff** \mathcal{A} is false on that variable assignment over that interpretation;
- › $\mathcal{A} \wedge \mathcal{B}$ is true on a variable assignment over an interpretation **iff** both \mathcal{A} is true on that variable assignment over that interpretation and \mathcal{B} is true on that variable assignment over that interpretation;
- › ...
- › $\forall x\mathcal{A}$ is true on a variable assignment over an interpretation **iff** \mathcal{A} is true on every variable assignment differing from the original one at most in what it assigns to x over that interpretation;
- › $\exists x\mathcal{A}$ is true on a variable assignment over an interpretation **iff** \mathcal{A} is true on some variable assignment differing from the original one at most in what it assigns to x over that interpretation.

The last two clauses are where this approach is strongest. Rather than considering some substitution instance of $\forall x\mathcal{A}$, we simply consider the direct constituent \mathcal{A} . Rather than considering all variations on the original interpretation which include some new name, we simply consider all ways of varying what is assigned to x by the original variable assignment, but keeping everything else unchanged. If you see the rationale for the clauses offered in §22.3, you can see why the clauses just offered are appropriate.

We now have the idea of truth on a variable assignment over an interpretation. But what we want – if this alternative approach is to yield the same end result – is truth in an interpretation. Notice that, given an interpretation, varying the variable assignment can change the truth value of a formula with free variables. But it cannot change the truth value of a formula which is a sentence, so if a sentence is true on one variable assignment over an interpretation, it is true on every variable assignment over that interpretation. So we can reintroduce the notion of truth in an interpretation, like so:

\mathcal{A} is true in an interpretation **iff**
 \mathcal{A} is a sentence of Quantifier, and \mathcal{A} is true on any (or every) variable assignment over that interpretation.

Here's how this works in practice. Suppose we want to figure out whether the sentence ' $\forall x\exists yLxy$ ' is true on an interpretation which associates the predicate ' Lxy ' with the relation ' x is no heavier than y ', and has as its domain the planets in our solar system. We might reason as follows:

For each way of picking planets to be the values of ' x ' and ' y ', either we pick two different planets, and one is lighter than the other (no two planets have the same mass); or we pick the same planet, and 'they' are identical in mass. In either

case, we can always find something no heavier than anything we pick. So for any variable assignment to 'x' over this interpretation, we can then assign something to 'y' so as to make 'Lxy' true on that joint assignment. Hence no matter what we assign to 'x', '∃yLxy' is true on that assignment. But since that is true no matter what we assign to 'x', '∀x∃yLxy' is true on every variable assignment over this interpretation. But that latter is a sentence, so is true in this interpretation.

Let us say that a sequence of objects $\langle a_1, \dots, a_n \rangle$ SATISFIES a formula \mathcal{A} in which the variables x_1, \dots, x_n occur freely iff there is a variable assignment over an interpretation whose domain includes each a_i , and which assigns each a_i to the variable x_i , and on which \mathcal{A} is true over that interpretation. What we have expressed in terms of variable assignments could have been expressed, a little more awkwardly, using the notion of some objects satisfying a formula. Indeed, this is how Alfred Tarski, the inventor of this approach to truth in Quantifier, first introduced the idea.²

Practice exercises

A. Consider the following interpretation:

- › The domain comprises only Corwin and Benedict
- › 'Ax' is to be true of both Corwin and Benedict
- › 'Bx' is to be true of Benedict only
- › 'Nx' is to be true of no one
- › 'c' is to refer to Corwin

Determine whether each of the following sentences is true or false in that interpretation:

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

B. Consider the following interpretation:

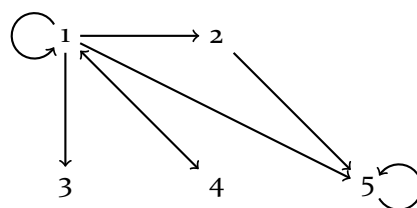
² A translation of his original 1933 paper is 'The Concept of Truth in Formalized Languages', 1983, in Tarski, *Logic, Semantics, Metamathematics*, Indianapolis: Hackett, pp. 152–278. It is rather technical in places.

- › The domain comprises only Lemmy, Courtney and Eddy
- › ‘ Gx ’ is to be true of Lemmy, Courtney and Eddy.
- › ‘ Hx ’ is to be true of and only of Courtney
- › ‘ Mx ’ is to be true of and only of Lemmy and Eddy
- › ‘ c ’ is to refer to Courtney
- › ‘ e ’ is to refer to Eddy

Determine whether each of the following sentences is true or false in that interpretation:

1. Hc
2. He
3. $Mc \vee Me$
4. $Gc \vee \neg Gc$
5. $Mc \rightarrow Gc$
6. $\exists x Hx$
7. $\forall x Hx$
8. $\exists x \neg Mx$
9. $\exists x (Hx \wedge Gx)$
10. $\exists x (Mx \wedge Gx)$
11. $\forall x (Hx \vee Mx)$
12. $\exists x Hx \wedge \exists x Mx$
13. $\forall x (Hx \leftrightarrow \neg Mx)$
14. $\exists x Gx \wedge \exists x \neg Gx$
15. $\forall x \exists y (Gx \wedge Hy)$

C. Following the diagram conventions introduced at the end of §21, consider the following interpretation:



Determine whether each of the following sentences is true or false in that interpretation:

1. $\exists x Rxx$
2. $\forall x Rxx$
3. $\exists x \forall y Rxy$
4. $\exists x \forall y Ryx$
5. $\forall x \forall y \forall z ((Rxy \wedge Ryz) \rightarrow Rxz)$
6. $\forall x \forall y \forall z ((Rxy \wedge Rxz) \rightarrow Ryz)$
7. $\exists x \forall y \neg Rxy$
8. $\forall x (\exists y Rxy \rightarrow \exists y Ryx)$
9. $\exists x \exists y (\neg x = y \wedge Rxy \wedge Ryx)$
10. $\exists x \forall y (Rxy \leftrightarrow x = y)$
11. $\exists x \forall y (Ryx \leftrightarrow x = y)$
12. $\exists x \exists y (\neg x = y \wedge Rxy \wedge \forall z (Rzx \leftrightarrow y = z))$

23 Semantic concepts

Offering a precise definition of truth in *Quantifier* was more than a little fiddly. But now that we are done, we can define various central logical notions. These will look very similar to the definitions we offered for *Sentential*. However, remember that they concern *interpretations*, rather than valuations.

We will use the symbol '⊨' for entailment in *Quantifier*, much as we did for *Sentential*. This introduces an ambiguity, but a harmless one, since every valid argument in *Sentential* remains a valid argument in *Quantifier*.

So:

- › $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ together ENTAIL \mathcal{C} (written $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \models \mathcal{C}$) iff there is no interpretation in which all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are true and in which \mathcal{C} is false.
- › An *Quantifier* sentence \mathcal{A} is a LOGICAL TRUTH iff \mathcal{A} is true in every interpretation; this is written $\models \mathcal{A}$.
- › \mathcal{A} is a CONTRADICTION iff \mathcal{A} is false in every interpretation; i.e., $\models \neg \mathcal{A}$.
- › Two *Quantifier* sentences \mathcal{A} and \mathcal{B} are LOGICALLY EQUIVALENT iff they are true in exactly the same interpretations as each other; i.e., both $\mathcal{A} \models \mathcal{B}$ and $\mathcal{B} \models \mathcal{A}$.
- › The *Quantifier* sentences $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are JOINTLY CONSISTENT iff there is some interpretation in which all of the sentences are true. They are JOINTLY INCONSISTENT iff there is no such interpretation.

As before, there is a close connection between validity and entailment:

$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$ is VALID in *Quantifier* iff there is no interpretation in which all of the premises are true and the conclusion is false; i.e., $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \models \mathcal{C}$. It is INVALID in *Quantifier* otherwise.

Our entailments are conclusive in virtue of form, because the consideration of many alternative interpretations of the non-logical expressions involved means that no argument

can be conclusive because of some special ‘connection in meaning’ between such expressions in the language. Entailment in *Quantifier* reflects only the logical structure of the sentences involved, where this now includes the new logical resources of *Quantifier* over *Sentential*.

23.1 Some subtleties about truth and interpretation

Our interpretations allow us a precise characterisation of formal truth for *Quantifier*. In §2, we introduced the idea of the structure of a sentence, and the pragmatic approach that logicians take to that notion. In *Quantifier*, the logical constants are those of *Sentential* plus the quantifiers and the identity predicate. But every other item of the language is up for reinterpretation, and our interpretations allow us to reinterpret these expressions in an arbitrary fashion.

In *Sentential*, we were able to characterise a logical truth as a sentence which is actually true on every reinterpretation of non-logical expressions (in *Sentential*, the non-logical expressions are just the atomic sentences). And we characterised a formally valid argument as one in which each possible reinterpretation (i.e., valuation) of the premises which makes them actually true, is also one which makes the conclusion actually true.

But this account of formal validity needs refining when it comes to *Quantifier*. Consider the *Quantifier* sentence which says that there are at least three things:

$$\exists x \exists y \exists z (x \neq y \wedge x \neq z \wedge y \neq z).$$

This sentence contains no non-logical expressions. Hence the sentence has a constant meaning, and is true on each reinterpretation just in case it is actually true. It is actually true – there are actually at least three things. Hence our *Sentential*-inspired account of logical truth would suggest that this sentence is a logical truth. But it seems very strange to think that it is a truth of logic alone that there are at least three things. Isn’t it possible that there had been fewer?

The keen-eyed among you will have noticed that this sentence is not, in fact, a logical truth of *Quantifier*. The reason is that in defining truth for sentences of *Quantifier*, we considered not just reinterpretations of the non-logical expressions, but also we allowed the domain of our interpretation to freely depart from actuality. So we are, in effect, allowing our interpretations to vary the meanings of the non-logical vocabulary *and* to vary the possible situations at which our sentences are to be evaluated. So we can consider a possible situation in which there are just two things, and if that situation provides the domain of an interpretation, there is no way of extending that interpretation to three names ‘*a*’, ‘*b*’, and ‘*c*’ such that ‘ $a \neq b \wedge a \neq c \wedge b \neq c$ ’ is true; hence ‘ $\exists x \exists y \exists z (x \neq y \wedge x \neq z \wedge y \neq z)$ ’ isn’t true on the original interpretation.

One sentence of this class is nevertheless a logical truth: the claim that something exists, ‘ $\exists x x = x$ ’. Since it is a constraint on *Quantifier* domains that they cannot be empty (§15.5), there is for any interpretation a way of adding a new name ‘*c*’ which denotes an

object in the domain, and of course ' $c = c$ ' is true in that extended interpretation, since each name must have a referent in any interpretation, and the identity predicate is always interpreted as representing pairs of objects in the domain and themselves.

24 Demonstrating Consistency and Invalidity

24.1 Logical truths and contradictions

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

In order for $\exists xAxx \rightarrow Bd$ to be false, the antecedent ($\exists xAxx$) must be true, and the consequent (Bd) must be false. To construct such an interpretation, we start by specifying a domain. Keeping the domain small makes it easier to specify what the predicates will be true of, so we shall start with a domain that has just one member. For concreteness, let's say it is the city of Paris.

domain: Paris

The name ' d ' must name something in the domain, so we have no option but:

d : Paris

Recall that we want $\exists xAxx$ to be true, so we want all members of the domain to be paired with themselves in the extension of ' A '. We can offer:

Axy : _____ x is in the same place as _____ y

Now ' Add ' is true, so it is surely true that $\exists xAxx$. Next, we want ' Bd ' to be false, so the referent of ' d ' must not be in the extension of ' B '. We might simply offer:

Bx : _____ x is in Germany

Now we have an interpretation where $\exists xAxx$ is true, but where ' Bd ' is false. So there is an interpretation where $\exists xAxx \rightarrow Bd$ is false. So $\exists xAxx \rightarrow Bd$ is not a logical truth.

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

domain: Paris

d : Paris

Axy : _____ x is in the same place as _____ y

Bx : _____ x is in France

This shows that there is an interpretation where ' $\exists xAxx \rightarrow Bd$ ' is true. So ' $\exists xAxx \rightarrow Bd$ ' is not a contradiction.

24.2 Logical equivalence

Suppose we want to show that ' $\forall xSx$ ' and ' $\exists xSx$ ' are not logically equivalent. We need to construct an interpretation 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 domain. Again, we make the domain small so that we can specify extensions easily. In this case, we shall need at least two objects. (If we chose a domain with only one member, the two sentences would end up with the same truth value. In order to see why, try constructing some partial interpretations with one-member domains.) For concreteness, let's take:

domain: Ornette Coleman, Sarah Vaughan

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 '. For concreteness we shall offer:

Sx : _____ x plays saxophone

Now ' $\exists xSx$ ' is true, because ' Sx ' is true of Ornette Coleman. Slightly more precisely, extend our interpretation by allowing ' c ' to name Ornette Coleman. ' Sc ' is true in this extended interpretation, so ' $\exists xSx$ ' was true in the original interpretation. Similarly, ' $\forall xSx$ ' is false, because ' Sx ' is false of Sarah Vaughan. Slightly more precisely, extend our interpretation by allowing ' d ' to name Sarah Vaughan, and ' Sd ' is false in this extended interpretation, so ' $\forall xSx$ ' was false in the original interpretation. We have provided a counter-interpretation to the claim that ' $\forall xSx$ ' and ' $\exists xSx$ ' are logically equivalent.

To show that \mathcal{A} is not a logical truth, it suffices to find an interpretation where \mathcal{A} is false.

To show that \mathcal{A} is not a contradiction, it suffices to find an interpretation where \mathcal{A} is true.

To show that \mathcal{A} and \mathcal{B} are not logically equivalent, it suffices to find an interpretation where one is true and the other is false.

24.3 Validity, entailment and consistency

To test for validity, entailment, or consistency, we typically need to produce interpretations that determine the truth value of several sentences simultaneously.

Consider the following argument in Quantifier:

$$\exists x(Gx \rightarrow Ga) \therefore \exists xGx \rightarrow Ga$$

To show that this is invalid, we must make the premise true and the conclusion false. The conclusion is a conditional, so to make it false, the antecedent must be true and the consequent must be false. Clearly, our domain must contain two objects. Let's try:

domain: Karl Marx, Ludwig von Mises

Gx : _____ x hated communism

a : Karl Marx

Given that Marx wrote *The Communist Manifesto*, ' Ga ' is plainly false in this interpretation. But von Mises famously hated communism. So ' $\exists xGx$ ' is true in this interpretation. Hence ' $\exists xGx \rightarrow Ga$ ' is false, as required.

But does this interpretation make the premise true? Yes it does! Note that ' $Ga \rightarrow Ga$ ' is true. (Indeed, it is a logical truth.) But then certainly ' $\exists x(Gx \rightarrow Ga)$ ' is true. So the premise is true, and the conclusion is false, in this interpretation. The argument is therefore invalid.

In passing, note that we have also shown that ' $\exists x(Gx \rightarrow Ga)$ ' does *not* entail ' $\exists xGx \rightarrow Ga$ '. And equally, we have shown that the sentences ' $\exists x(Gx \rightarrow Ga)$ ' and ' $\neg(\exists xGx \rightarrow Ga)$ ' are jointly consistent.

Let's consider a second example. Consider:

$$\forall x\exists yLxy \therefore \exists y\forall xLxy$$

Again, I want to show that this is invalid. To do this, we must make the premises true and the conclusion false. Here is a suggestion:

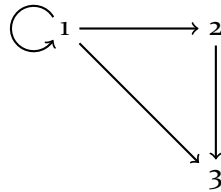
domain: People with a living biological sibling

Lxy : _____ x shares a parent with _____ y

The premise is clearly true on this interpretation. Anyone in the domain has a living sibling. That sibling will also, then, be in the domain, because one cannot be someone's sibling without also having them as a sibling. So for everyone in the domain, there will be at least one other person in the domain who is their sibling, and thus has a parent in common with them. Hence ' $\forall x\exists yLxy$ ' is true. But the conclusion is clearly false, for that would require that there is some single person who shares a parent with everyone in the domain, and there is no such person. So the argument is invalid. We observe immediately

that the sentences ‘ $\forall x\exists yLxy$ ’ and ‘ $\neg\exists y\forall xLxy$ ’ are jointly consistent and that ‘ $\forall x\exists yLxy$ ’ does not entail ‘ $\exists y\forall xLxy$ ’.

For my third example, I shall mix things up a bit. In §21, I described how we can present some interpretations using diagrams. For example:



Using the conventions employed in §21, the domain of this interpretation is the first three positive whole numbers, and ‘ Rxy ’ is true of x and y just in case there is an arrow from x to y in our diagram. Here are some sentences that the interpretation makes true:

- > ‘ $\forall x\exists yRyx$ ’
- > ‘ $\exists x\forall yRxy$ ’ witness 1
- > ‘ $\exists x\forall y(Ryx \leftrightarrow x = y)$ ’ witness 1
- > ‘ $\exists x\exists y\exists z(\neg y = z \wedge Rxy \wedge Rzx)$ ’ witness 2
- > ‘ $\exists x\forall y\neg Rxy$ ’ witness 3
- > ‘ $\exists x(\exists yRyx \wedge \neg\exists yRxy)$ ’ witness 3

This immediately shows that all of the preceding six sentences are jointly consistent. We can use this observation to generate *invalid* arguments, e.g.:

$$\forall x\exists yRyx, \exists x\forall yRxy \therefore \forall x\exists yRxy$$

$$\exists x\forall yRxy, \exists x\forall y\neg Rxy \therefore \neg\exists x\exists y\exists z(\neg y = z \wedge Rxy \wedge Rzx)$$

and many more besides.

To show that $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$ is invalid, it suffices to find an interpretation where all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are true and where \mathcal{C} is false.
 That same interpretation will show that $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ do not entail \mathcal{C} .
 That same interpretation will show that $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n, \neg\mathcal{C}$ are jointly consistent.

When you provide an interpretation to refute a claim – to logical truth, say, or to entailment – this is sometimes called providing a *counter-interpretation* (or providing a *counter-model*).

Practice exercises

A. Show that each of the following is neither a logical truth nor a contradiction:

1. $Da \wedge Db$
2. $\exists xTxh$
3. $Pm \wedge \neg\forall xPx$
4. $\forall zJz \leftrightarrow \exists yJy$
5. $\forall x(Wxmn \vee \exists yLxy)$
6. $\exists x(Gx \rightarrow \forall yMy)$
7. $\exists x(x = h \wedge x = i)$

B. 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. $\forall x(Px \rightarrow \neg Qx), \exists x(Px \wedge \neg Qx)$
6. $\exists x(Px \wedge Qx), \exists x(Px \rightarrow Qx)$
7. $\forall x(Px \rightarrow Qx), \forall x(Px \wedge Qx)$
8. $\forall x\exists yRxy, \exists x\forall yRxy$
9. $\forall x\exists yRxy, \forall x\exists yRyx$

C. Show that the following sentences are jointly consistent:

1. $Ma, \neg Na, Pa, \neg Qa$
2. $Lee, Leg, \neg Lge, \neg Lgg$
3. $\neg(Ma \wedge \exists xAx), Ma \vee Fa, \forall x(Fx \rightarrow Ax)$
4. $Ma \vee Mb, Ma \rightarrow \forall x\neg Mx$
5. $\forall yGy, \forall x(Gx \rightarrow Hx), \exists y\neg Iy$
6. $\exists x(Bx \vee Ax), \forall x\neg Cx, \forall x((Ax \wedge Bx) \rightarrow Cx)$
7. $\exists xXx, \exists xYx, \forall x(Xx \leftrightarrow \neg Yx)$
8. $\forall x(Px \vee Qx), \exists x\neg(Qx \wedge Px)$
9. $\exists z(Nz \wedge Ozz), \forall x\forall y(Oxy \rightarrow Oyx)$
10. $\neg\exists x\forall yRxy, \forall x\exists yRxy$
11. $\neg Raa, \forall x(x = a \vee Rxa)$
12. $\forall x\forall y\forall z(x = y \vee y = z \vee x = z), \exists x\exists y \neg x = y$
13. $\exists x\exists y(Zx \wedge Zy \wedge x = y), \neg Zd, d = e$

D. Show that the following arguments are invalid:

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

25 Reasoning about all interpretations: demonstrating inconsistency and validity

25.1 Logical truths and contradictions

We can show that a sentence is *not* a logical truth just by providing one carefully specified interpretation: an interpretation in which the sentence is false. To show that something is a logical truth, on the other hand, it would not be enough to construct ten, one hundred, or even a thousand interpretations in which the sentence is true. A sentence is only a logical truth if it is true in *every* interpretation, and there are infinitely many interpretations. We need to reason about all of them, and we cannot do this by dealing with them one by one!

Sometimes, we can reason about all interpretations fairly easily. For example, we can offer a relatively simple argument that ' $Raa \leftrightarrow Raa$ ' is a logical truth:

Any relevant interpretation will give ' Raa ' a truth value. If ' Raa ' is true in an interpretation, then ' $Raa \leftrightarrow Raa$ ' is true in that interpretation. If ' Raa ' is false in an interpretation, then ' $Raa \leftrightarrow Raa$ ' is true in that interpretation. These are the only alternatives. So ' $Raa \leftrightarrow Raa$ ' is true in every interpretation. Therefore, it is a logical truth.

This argument is valid, of course, and its conclusion is true. However, it is not an argument in *Quantifier*. Rather, it is an argument in English *about* *Quantifier*: it is an argument in the metalanguage.

Note another feature of the argument. Since the sentence in question contained no quantifiers, we did not need to think about how to interpret ' a ' and ' R '; the point was just that, however we interpreted them, ' Raa ' would have some truth value or other. (We could ultimately have given the same argument concerning *Sentential* sentences.)

Here is another bit of reasoning. Consider the sentence ' $\forall x(Rxx \leftrightarrow Rxx)$ '. Again, it should obviously be a logical truth. But to say precisely why is quite a challenge. We cannot say that ' $Rxx \leftrightarrow Rxx$ ' is true in every interpretation, since ' $Rxx \leftrightarrow Rxx$ ' is not even a *sentence* of *Quantifier* (remember that ' x ' is a variable, not a name). So we have to be a bit cleverer.

Consider some arbitrary interpretation. Consider some arbitrary member of the model's domain, which, for convenience, we shall call *obbie*, and suppose we extend our original interpretation by adding a new name, 'c', to name *obbie*. Then either '*Rcc*' will be true or it will be false. If '*Rcc*' is true, then ' $Rcc \leftrightarrow Rcc$ ' is true. If '*Rcc*' is false, then ' $Rcc \leftrightarrow Rcc$ ' will be true. So either way, ' $Rcc \leftrightarrow Rcc$ ' is true. Since there was nothing special about *obbie* – we might have chosen any object – we see that no matter how we extend our original interpretation by allowing 'c' to name some new object, ' $Rcc \leftrightarrow Rcc$ ' will be true in the new interpretation. So ' $\forall x(Rxx \leftrightarrow Rxx)$ ' was true in the original interpretation. But we chose our interpretation arbitrarily. So ' $\forall x(Rxx \leftrightarrow Rxx)$ ' is true in every interpretation. It is therefore a logical truth.

This is quite longwinded, but, as things stand, there is no alternative. In order to show that a sentence is a logical truth, we must reason about *all* interpretations.

But sometimes we can draw a conclusion about some way all interpretations must be, by considering a hypothetical interpretation which isn't that way, and showing that hypothesis leads to trouble. Here is an example.

Suppose there were an interpretation which makes ' $a \neq b \rightarrow \neg \exists x(x = a \wedge x = b)$ ' false. Then while ' $a \neq b$ ' is true on that interpretation, ' $\neg \exists x(x = a \wedge x = b)$ ' must be false. So ' $\exists x(x = a \wedge x = b)$ ' must be true; and hence it has some true substitution instance with some new name 'c' as the instantiating name: ' $c = a \wedge c = b$ ' is true. But then it must also be that ' $a = b$ ' is true, which contradicts our original supposition that ' $a \neq b$ ' is true on this interpretation. So there can be no such interpretation, and hence the sentence cannot be false on any interpretation and must be a logical truth.

This is an example of *REDUCTIO* reasoning: we assume something, and derive some absurdity or contradiction; we then conclude that our assumption is to be rejected. It is often easier to show that all Fs have a property by deriving an absurdity from the assumption that some F lacks the property, than it is to show it 'directly' by considering each F in turn. This is a powerful technique, used widely in mathematics and elsewhere.

25.2 Other cases

Similar points hold of other cases too. Thus, we must reason about all interpretations if we want to show:

- › that a sentence is a contradiction; for this requires that it is false in *every* interpretation.
- › that two sentences are logically equivalent; for this requires that they have the same truth value in *every* interpretation.

- › that some sentences are jointly inconsistent; for this requires that there is no interpretation in which all of those sentences are true together; i.e., that, in *every* interpretation, at least one of those sentences is false.
- › that an argument is valid; for this requires that the conclusion is true in *every* interpretation where the premises are true.
- › that some sentences entail another sentence.

The problem is that, with the tools available to you so far, reasoning about all interpretations is a serious challenge! Let's take just one more example. Here is an argument which is obviously valid:

$$\forall x(Hx \wedge Jx) \therefore \forall xHx$$

After all, if everything is both H and J, then everything is H. But we can only show that the argument is valid by considering what must be true in every interpretation in which the premise is true. And to show this, we would have to reason as follows:

Consider an arbitrary interpretation in which the premise ' $\forall x(Hx \wedge Jx)$ ' is true. It follows that, however we expand the interpretation with a new name, for example ' c ', ' $Hc \wedge Jc$ ' will be true in this new interpretation. ' Hc ' will, then, also be true in this new interpretation. But since this held for *any* way of expanding the interpretation, it must be that ' $\forall xHx$ ' is true in the old interpretation. And we assumed nothing about the interpretation except that it was one in which ' $\forall x(Hx \wedge Jx)$ ' is true. So any interpretation in which ' $\forall x(Hx \wedge Jx)$ ' is true is one in which ' $\forall xHx$ ' is true. The argument is valid!

Even for a simple argument like this one, the reasoning is somewhat complicated. For longer arguments, the reasoning can be extremely torturous.

Reductio reasoning is particularly useful in demonstrating the validity of valid arguments. In that case we will typically make the hypothetical supposition that there is some interpretation which makes the premises true and the conclusion false, and then derive an absurdity from that supposition. From that we can conclude there is no such interpretation, and hence that any interpretation which makes the premises true must also be one which makes the conclusion true.

The following table summarises whether a single (counter-)interpretation suffices, or whether we must reason about all interpretations (whether directly considering them all, or indirectly making use of *reductio*).

	Yes	No
logical truth?	all interpretations	one counter-interpretation
contradiction?	all interpretations	one counter-interpretation
equivalent?	all interpretations	one counter-interpretation
consistent?	one interpretation	consider all interpretations
valid?	all interpretations	one counter-interpretation
entailment?	all interpretations	one counter-interpretation

This might usefully be compared with the table at the end of §13. The key difference resides in the fact that *Sentential* concerns truth tables, whereas *Quantifier* concerns interpretations. This difference is deeply important, since each truth-table only ever has finitely many lines, so that a complete truth table is a relatively tractable object. By contrast, there are infinitely many interpretations for any given sentence(s), so that reasoning about all interpretations can be a deeply tricky business.

Chapter 6

Natural deduction for Sentential

26 The very idea of natural deduction

Way back in §2, we said that an argument is valid iff it is impossible to make all of the premises true and the conclusion false.

In the case of *Sentential*, this led us to develop truth tables. Each line of a complete truth table corresponds to a valuation. So, when faced with a *Sentential* argument, we have a very direct way to assess whether it is possible to make all of the premises true and the conclusion false: just thrash through the truth table.

But truth tables do not necessarily give us much *insight*. Consider two arguments in *Sentential*:

$$P \vee Q, \neg P \therefore Q$$

$$P \leftrightarrow Q, P \therefore Q.$$

Clearly, these are valid arguments. You can confirm that they are valid by constructing four-line truth tables. But we might say that two arguments, proceeding from different premises with different logical connectives involves, must make use of different principles of implication. What follows from a disjunction is not at all the same as what follows from a biconditional, And it might be nice to keep track of these differences.

Certainly human reasoning treats disjunctions very differently from biconditionals. While logic is not really about human reasoning, which is more properly a subject matter for psychology, nevertheless we can formally study the different forms of reasoning involved in arguing from sentences with different structures, by asking: *what would it be acceptable to conclude from a premise with a certain formal structure, supposing one cannot give up one's commitment to that premise?*

One aim of a *natural deduction system* is to show that particular arguments are valid, in a way that allows us to understand the reasoning that the arguments might involve in virtue of the forms of the claims involved. We begin with very basic rules of implication. These rules can be combined, to offer more complicated arguments. Indeed, with just a small starter pack of rules of implication, we hope to capture all valid arguments.

This is a very different way of thinking about arguments.

With truth tables, we directly consider different ways to make sentences true or false. With natural deduction systems, we manipulate sentences in accordance with rules that

we have set down as good rules. The latter promises to give us a better insight – or at least, a different insight – into how arguments work. Consider this pair of arguments:

$$\begin{aligned} &\neg P \vee Q, P \therefore Q \\ &P \rightarrow Q, P \therefore Q. \end{aligned}$$

From a truth-table perspective, these arguments are indistinguishable: their premises are true in the same valuations, and so are their conclusions. Yet they are distinct from an implicational point of view, since one will involve consideration of the consequences of disjunctions and negations, and the other consideration of the consequences of the conditional. So the natural deduction perspective on arguments promises to give us a more fine-grained analysis of how valid arguments function.

The move to natural deduction might be motivated by more than the search for insight. It might also be motivated by *practicality*. Consider:

$$A_1 \rightarrow C_1 \therefore (A_1 \wedge A_2 \wedge A_3 \wedge A_4 \wedge A_5) \rightarrow (C_1 \vee C_2 \vee C_3 \vee C_4 \vee C_5).$$

To test this argument for validity, you might use a 1024-line truth table. If you do it correctly, then you will see that there is no line on which all the premises are true and on which the conclusion is false. So you will know that the argument is valid. (But, as just mentioned, there is a sense in which you will not know *why* the argument is valid.) But now consider:

$$\begin{aligned} A_1 \rightarrow C_1 \therefore (A_1 \wedge A_2 \wedge A_3 \wedge A_4 \wedge A_5 \wedge A_6 \wedge A_7 \wedge A_8 \wedge A_9 \wedge A_{10}) \rightarrow \\ (C_1 \vee C_2 \vee C_3 \vee C_4 \vee C_5 \vee C_6 \vee C_7 \vee C_8 \vee C_9 \vee C_{10}). \end{aligned}$$

This argument is also valid – as you can probably tell – but to test it requires a truth table with $2^{20} = 1048576$ lines. In principle, we can set a machine to grind through truth tables and report back when it is finished. In practice, complicated arguments in Sentential can become *intractable* if we use truth tables.

When we get to Quantifier, though, the problem gets dramatically worse. There is nothing like the truth table test for Quantifier. To assess whether or not an argument is valid, we have to reason about *all* interpretations. But there are infinitely many possible interpretations. We cannot even in principle set a machine to grind through infinitely many possible interpretations and report back when it is finished: it will *never* finish. We either need to come up with some more efficient way of reasoning about all interpretations, or we need to look for something different. Since we already have some motivation for considering the role in arguments of particular premises, rather than all the premises collectively as in the truth-table approach, we will opt for the ‘something different’ path – natural deduction.

Rather than reasoning directly about all valuations (in the case of Sentential) or all interpretations (in the case of Quantifier), we shall try to select a few basic rules of implication. Some of these will govern the behaviour of the sentential connectives. Others will

govern the behaviour of the quantifiers and identity. The resulting system of rules will give us a new way to think about the validity of arguments. The modern development of natural deduction dates from simultaneous and unrelated papers by Gerhard Gentzen and Stanisław Jaśkowski (1934). However, the natural deduction system that we shall consider is based largely around work by Frederic Fitch (first published in 1952).

Natural deduction was so-called because the rules of implication it codifies were seen as reflecting 'natural' forms of human reasoning. It must be admitted that no one spontaneously and instinctively reasons in a way that conforms to the rules of natural deduction. But there is one place where these forms of inference are widespread – mathematical reasoning. And it will not surprise you to learn that these systems of inference were introduced initially to codify good practice in mathematical proofs. Don't worry, though: we won't expect that you are already a fluent mathematician. Though some of the rules might be a bit stilted and formal for everyday use, the rationale for each of them is transparent and can be easily understood even by those without extensive mathematical training.

27 Basic rules for Sentential

We will develop a NATURAL DEDUCTION system. For each connective, there will be INTRODUCTION rules, that allow us to prove a sentence that has that connective as the main connective, and ELIMINATION rules, that allow us to prove something given a sentence that has that connective as the main connective

27.1 The idea of a formal proof

Very abstractly, a FORMAL PROOF in a natural deduction system is a sequence of sentences, each accompanied by a ‘commentary’ on that sentence. Some of those sentences may be initial *assumptions*, or suppositions ‘for the sake of argument’. Other sentences will follow from earlier sentences in the formal proof by following the natural deduction rules we will detail below. The commentary tells us whether the sentence is an assumption, or whether it follows from one of the rules from an earlier sentence or sentences, and specifies which sentence(s) it follows from, and which assumptions it relies on. The last line of the formal proof is the conclusion. Some of the initial assumptions may remain as assumptions or suppositions even at the end of the formal proof – these will be the premises of an argument that the formal proof demonstrates to be valid. Henceforth, I shall simply call these ‘proofs’, but you should be aware that there are *informal proofs* too.¹

We will use a particular graphical representation of natural deduction proofs, one which makes use of graphical nesting of sentences to vividly represent which assumptions a particular sentence in a proof is relying on at any given stage, and uses a device of horizontal marks to distinguish assumptions from derived sentences. It will be easier to see how this works by considering some examples!

So consider:

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

We shall start a proof by writing an assumption:

¹ Many of the arguments we offer in our metalanguage, quasi-mathematical arguments *about* our formal languages, are proofs. Sometimes people call formal proofs ‘deductions’ or ‘derivations’, to ensure that no one will confuse the metalanguage activity of proving things about our logical languages with the activity of constructing arguments within those languages. But it seems unlikely that anyone in this course will be confused on this point, since we are not offering very many proofs in the metalanguage in the first place!

$$1 \quad \underline{\neg(A \vee B)}$$

Note that we have numbered the premise, since we shall want to refer back to it. Indeed, every line on a proof is numbered, so that we can refer back to it.

Note also that we have drawn a line underneath the premise. Everything written above the line is an *assumption*. Everything written below the line will either be something which follows from the assumptions, or it will be some new assumption.

The vertical line at the side marks the range of this assumption: any subsequent sentence which is adjacent to a continuation of this vertical line is understood to lie in the *range* of this assumption. If we are going to write down some sentence that follows from our assumption, we will involve first extending that vertical line to indicate that it rests on this assumption, or at least is playing a role in the argument connected with this assumption. (As we will see, not every sentence in the range of an assumption will *depend* on the assumption, but most will.)

We are hoping to conclude that ' $\neg A \wedge \neg B$ '; so we are hoping ultimately to conclude our proof with

$$n \quad \neg A \wedge \neg B$$

for some number n . It doesn't matter which line we end on, but we would obviously prefer a short proof to a long one.

Similarly, suppose we wanted to consider:

$$A \vee B, \neg(A \wedge C), \neg(B \wedge \neg D) \therefore \neg C \vee D.$$

The argument has three premises, so we start by writing them all down, numbered, and drawing a line under them:

$$\begin{array}{l} 1 \\ 2 \\ 3 \end{array} \quad \underline{\begin{array}{l} A \vee B \\ \neg(A \wedge C) \\ \neg(B \wedge \neg D) \end{array}}$$

and we are hoping to conclude with some line:

$$n \quad \neg C \vee D$$

All that remains to do is to explain each of the rules that we can use along the way from premises to conclusion. The rules are broken down by our logical connectives.

27.2 Conjunction

Suppose I want to show that Ludwig is reactionary *and* libertarian. One obvious way to do this would be as follows: first I show that Ludwig is reactionary; then I show that Ludwig is libertarian; then I put these two demonstrations together, to obtain the conjunction.

Our natural deduction system will capture this thought straightforwardly. In the example given, I might adopt the following symbolisation key:

R : Ludwig is reactionary
 L : Ludwig is libertarian

Perhaps I am working through a proof, and I have obtained ' R ' on line 8 and ' L ' on line 15. Then on any subsequent line I can obtain ' $R \wedge L$ ' thus:

$$\begin{array}{l|l} 8 & R \\ 15 & L \\ & R \wedge L \quad \wedge I 8, 15 \end{array}$$

Note that every line of our proof must either be an assumption, or must be justified by some rule. We cite ' $\wedge I 8, 15$ ' here to indicate that the line is obtained by the rule of conjunction introduction ($\wedge I$) applied to lines 8 and 15. This is our way of representing what I earlier called the *commentary* on this sentence as it appears in the proof. The vertical and horizontal lines are also parts of this commentary.

I could equally well have obtained:

$$\begin{array}{l|l} 8 & R \\ 15 & L \\ & L \wedge R \quad \wedge I 15, 8 \end{array}$$

with the citation reversed, to reflect the order of the conjuncts. More generally, here is our CONJUNCTION INTRODUCTION rule:

$\begin{array}{l l} m & \mathcal{A} \\ & \vdots \\ n & \mathcal{B} \\ & \vdots \\ & \mathcal{A} \wedge \mathcal{B} \quad \wedge I m, n \end{array}$

To be clear, the statement of the rule is *schematic*. It is not itself a proof. ‘ \mathcal{A} ’ and ‘ \mathcal{B} ’ are not sentences of Sentential. Rather, they are symbols in the metalanguage, which we use when we want to talk about any sentence of Sentential (see §7). Similarly, ‘ m ’ and ‘ n ’ are not numerals that will appear on any actual proof. Rather, they are symbols in the metalanguage, which we use when we want to talk about any line number of any proof. In an actual proof, the lines are numbered ‘1’, ‘2’, ‘3’, and so forth. But when we define the rule, we use variables to emphasise that the rule may be applied at any point. The rule requires only that we have both conjuncts available to us somewhere in the proof, earlier than the line that results from the application of the rule. They can be separated from one another, and they can appear in any order. So m might be less than n , or greater than n . Indeed, m might even equal n , as in this proof:

1	P	
2	$P \wedge P$	$\wedge I$ 1, 1

The rule is called ‘conjunction *introduction*’ because it introduces the symbol ‘ \wedge ’ into our proof where it may have been absent. Correspondingly, we have a rule that *eliminates* that symbol. Suppose you have shown that Ludwig is both reactionary and libertarian. You are entitled to conclude that Ludwig is reactionary. Equally, you are entitled to conclude that Ludwig is libertarian. Putting this together, we obtain our CONJUNCTION ELIMINATION rules:

<table style="border-collapse: collapse; margin: auto;"> <tr> <td style="padding-right: 10px;">m</td> <td style="border-left: 1px solid black; padding-left: 10px;">$\mathcal{A} \wedge \mathcal{B}$</td> <td></td> </tr> <tr> <td></td> <td style="border-left: 1px solid black; padding-left: 10px;">\vdots</td> <td></td> </tr> <tr> <td></td> <td style="border-left: 1px solid black; padding-left: 10px;">\mathcal{A}</td> <td style="padding-left: 20px;">$\wedge E$ m</td> </tr> </table>	m	$\mathcal{A} \wedge \mathcal{B}$			\vdots			\mathcal{A}	$\wedge E$ m	<table style="border-collapse: collapse; margin: auto;"> <tr> <td style="padding-right: 10px;">m</td> <td style="border-left: 1px solid black; padding-left: 10px;">$\mathcal{A} \wedge \mathcal{B}$</td> <td></td> </tr> <tr> <td></td> <td style="border-left: 1px solid black; padding-left: 10px;">\vdots</td> <td></td> </tr> <tr> <td></td> <td style="border-left: 1px solid black; padding-left: 10px;">\mathcal{B}</td> <td style="padding-left: 20px;">$\wedge E$ m</td> </tr> </table>	m	$\mathcal{A} \wedge \mathcal{B}$			\vdots			\mathcal{B}	$\wedge E$ m
m	$\mathcal{A} \wedge \mathcal{B}$																		
	\vdots																		
	\mathcal{A}	$\wedge E$ m																	
m	$\mathcal{A} \wedge \mathcal{B}$																		
	\vdots																		
	\mathcal{B}	$\wedge E$ m																	

The point is simply that, when you have a conjunction on some line of a proof, you can obtain either of the conjuncts by $\wedge E$. (One point, might be worth emphasising: you can only apply this rule when conjunction is the main connective. So you cannot derive ‘ D ’ just from ‘ $C \vee (D \wedge E)$ ’!)

Even with just these two rules, we can start to see some of the power of our formal proof system. Consider:

$$\begin{aligned} & ((A \vee B) \rightarrow (C \vee D)) \wedge ((E \vee F) \rightarrow (G \vee H)) \\ \therefore & ((E \vee F) \rightarrow (G \vee H)) \wedge ((A \vee B) \rightarrow (C \vee D)) \end{aligned}$$

The main connective in both the premise and conclusion of this argument is ‘ \wedge ’. In order to provide a proof, we begin by writing down the premise, which is our assumption. We draw a line below this: everything after this line must follow from our assumptions by (successive applications of) our rules of implication. So the beginning of the proof looks like this:

$$1 \quad \underline{[(A \vee B) \rightarrow (C \vee D)] \wedge [(E \vee F) \rightarrow (G \vee H)]}$$

From the premise, we can get each of the conjuncts by $\wedge E$. The proof now looks like this:

$$\begin{array}{l|l} 1 & \underline{[(A \vee B) \rightarrow (C \vee D)] \wedge [(E \vee F) \rightarrow (G \vee H)]} \\ 2 & [(A \vee B) \rightarrow (C \vee D)] \quad \wedge E 1 \\ 3 & [(E \vee F) \rightarrow (G \vee H)] \quad \wedge E 1 \end{array}$$

So by applying the $\wedge I$ rule to lines 3 and 2 (in that order), we arrive at the desired conclusion. The finished proof looks like this:

$$\begin{array}{l|l} 1 & \underline{[(A \vee B) \rightarrow (C \vee D)] \wedge [(E \vee F) \rightarrow (G \vee H)]} \\ 2 & [(A \vee B) \rightarrow (C \vee D)] \quad \wedge E 1 \\ 3 & [(E \vee F) \rightarrow (G \vee H)] \quad \wedge E 1 \\ 4 & [(E \vee F) \rightarrow (G \vee H)] \wedge [(A \vee B) \rightarrow (C \vee D)] \quad \wedge I 3, 2 \end{array}$$

This is a very simple proof, but it shows how we can chain rules of proof together into longer proofs. In passing, note that investigating this argument with a truth table would have required a staggering 256 lines; our formal proof required only four lines.

It is worth giving another example. Way back in §10.3, we noted that this argument is valid:

$$A \wedge (B \wedge C) \therefore (A \wedge B) \wedge C.$$

To provide a proof corresponding to this argument, we start by writing:

$$1 \quad \underline{A \wedge (B \wedge C)}$$

From the premise, we can get each of the conjuncts by applying $\wedge E$ twice. And we can then apply $\wedge E$ twice more, so our proof looks like:

$$\begin{array}{l|l} 1 & \underline{A \wedge (B \wedge C)} \\ 2 & A \quad \wedge E 1 \\ 3 & B \wedge C \quad \wedge E 1 \\ 4 & B \quad \wedge E 3 \\ 5 & C \quad \wedge E 3 \end{array}$$

But now we can merrily reintroduce conjunctions in the order we wanted them, so that our final proof is:

1	$A \wedge (B \wedge C)$	
2	A	$\wedge E$ 1
3	$B \wedge C$	$\wedge E$ 1
4	B	$\wedge E$ 3
5	C	$\wedge E$ 3
6	$A \wedge B$	$\wedge I$ 2, 4
7	$(A \wedge B) \wedge C$	$\wedge I$ 6, 5

Recall that our official definition of sentences in *Sentential* only allowed conjunctions with two conjuncts. When we discussed semantics, we became a bit more relaxed, and allowed ourselves to drop inner brackets in long conjunctions, since the order of the brackets did not affect the truth table. The proof just given suggests that we could also drop inner brackets in all of our proofs. However, this is not standard, and we shall not do this. Instead, we shall return to the more austere bracketing conventions. (Though we will allow ourselves to drop outermost brackets, for legibility.)

Our conjunction rules correspond to intuitively correct patterns of implication. But they are also demonstrably good in another sense, for our rules can be vindicated by considering facts about entailment. Each of these is easily demonstrated:

- › $\mathcal{A}, \mathcal{B} \models \mathcal{A} \wedge \mathcal{B}$;
- › $\mathcal{A} \wedge \mathcal{B} \models \mathcal{A}$;
- › $\mathcal{A} \wedge \mathcal{B} \models \mathcal{B}$.

It can be recognised, then, that our proof rules correspond to valid arguments in *Sentential*, and so our conjunction rules will never permit us to derive a false sentence from true sentences. (There is no guarantee of course that the assumptions we make in our formal proofs are in fact true, only that if they were true, what we derive from them would also be true.)

27.3 Conditional

Consider the following argument:

If Jane is smart then she is fast. Jane is smart. So Jane is fast.

This argument is certainly valid. And it suggests a straightforward CONDITIONAL ELIMINATION rule (\rightarrow E):

m	$\mathcal{A} \rightarrow \mathcal{B}$	
	\vdots	
n	\mathcal{A}	
	\vdots	
	\mathcal{B}	\rightarrow E m, n

This rule is also sometimes called *modus ponens*. Again, this is an elimination rule, because it allows us to obtain a sentence that may not contain ' \rightarrow ', having started with a sentence that did contain ' \rightarrow '. Note that the conditional, and the antecedent, can be separated from one another, and they can appear in any order. However, in the citation for \rightarrow E, we always cite the conditional first, followed by the antecedent.

The rule for conditional introduction is also quite easy to motivate. The following argument should be valid:

Ludwig is reactionary. Therefore if Ludwig is libertarian, then Ludwig is both reactionary *and* libertarian.

If someone doubted that this was valid, we might try to convince them otherwise by explaining ourselves as follows:

Assume that Ludwig is reactionary. Now, *additionally* assume that Ludwig is libertarian. Then by conjunction introduction – which we just discussed – Ludwig is both reactionary and libertarian. Of course, that's conditional on the assumption that Ludwig is libertarian. But this just means that, if Ludwig is libertarian, then he is both reactionary and libertarian.

Transferred into natural deduction format, here is the pattern of reasoning that we just used. We started with one premise, 'Ludwig is reactionary', thus:

1 R

The next thing we did is to make an *additional* assumption ('Ludwig is libertarian'), for the sake of argument. To indicate that we are no longer dealing *merely* with our original assumption ('R'), but with some additional assumption, we continue our proof as follows:

$$\begin{array}{l|l}
 1 & R \\
 \hline
 2 & \boxed{L}
 \end{array}$$

Note that we are *not* claiming, on line 2, to have proved ‘ L ’ from line 1. So we do not need to write in any justification for the additional assumption on line 2. We do, however, need to mark that it is an additional assumption. We do this by drawing a line under it (to indicate that it is an assumption) and by indenting it with a further vertical line (to indicate that it is additional).

With this extra assumption in place, we are in a position to use \wedge I. So we could continue our proof:

$$\begin{array}{l|l|l}
 1 & R & \\
 \hline
 2 & \boxed{L} & \\
 \hline
 3 & \boxed{R \wedge L} & \wedge I 1, 2
 \end{array}$$

So we have now shown that, on the additional assumption, ‘ L ’, we can obtain ‘ $R \wedge L$ ’. We can therefore conclude that, if ‘ L ’ obtains, then so does ‘ $R \wedge L$ ’. Or, to put it more briefly, we can conclude ‘ $L \rightarrow (R \wedge L)$ ’:

$$\begin{array}{l|l|l}
 1 & R & \\
 \hline
 2 & \boxed{L} & \\
 \hline
 3 & \boxed{R \wedge L} & \wedge I 1, 2 \\
 4 & L \rightarrow (R \wedge L) & \rightarrow I 2-3
 \end{array}$$

Observe that we have dropped back to using one vertical line. We have *discharged* the additional assumption, ‘ L ’, since the conditional itself follows just from our original assumption, ‘ R ’. The final line of this proof relies only on the undischarged assumption ‘ R ’ – we made use of the assumption ‘ L ’ only in the nested subproof, and the range of that assumption is restricted to sentences in that subproof.

The general pattern at work here is the following. We first make an additional assumption, A ; and from that additional assumption, we prove B . In that case, we have established the following: If it does in fact turn out that A , then it also turns out that B . This is wrapped up in the rule for **CONDITIONAL INTRODUCTION**:

i	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">\mathcal{A}</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">\vdots</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">j</td> <td style="padding: 0 5px;">\mathcal{B}</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">\vdots</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">$\mathcal{A} \rightarrow \mathcal{B}$</td> <td style="padding-left: 10px;">$\rightarrow I\ i-j$</td> </tr> </table>		\mathcal{A}			\vdots		j	\mathcal{B}			\vdots			$\mathcal{A} \rightarrow \mathcal{B}$	$\rightarrow I\ i-j$
	\mathcal{A}															
	\vdots															
j	\mathcal{B}															
	\vdots															
	$\mathcal{A} \rightarrow \mathcal{B}$	$\rightarrow I\ i-j$														

There can be as many or as few lines as you like between lines i and j . Notice that in our presentation of the rule, discharging the assumption \mathcal{A} takes us out of the subproof in which \mathcal{B} is derived from \mathcal{A} . If \mathcal{A} is the initial assumption of a proof, then discharging it may well leave us with a conditional claim that depends on no undischarged assumptions at all. We see an example in this proof, where the main proof, marked by the leftmost vertical line, features no horizontal line marking an assumption:

1	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">$P \wedge P$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">2</td> <td style="padding: 0 5px;">P</td> <td style="padding-left: 10px;">$\rightarrow E\ 1$</td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">3</td> <td style="padding: 0 5px;">$(P \wedge P) \rightarrow P$</td> <td style="padding-left: 10px;">$\rightarrow I\ 1-2$</td> </tr> </table>		$P \wedge P$		2	P	$\rightarrow E\ 1$	3	$(P \wedge P) \rightarrow P$	$\rightarrow I\ 1-2$
	$P \wedge P$									
2	P	$\rightarrow E\ 1$								
3	$(P \wedge P) \rightarrow P$	$\rightarrow I\ 1-2$								

It might come as no surprise that the conclusion of this proof – being provable from no undischarged assumptions at all – turns out to be a tautology.

It will help to offer a final illustration of $\rightarrow I$ in action. Suppose we want to consider the following:

$$P \rightarrow Q, Q \rightarrow R \therefore P \rightarrow R.$$

We start by listing *both* of our premises. Then, since we want to arrive at a conditional (namely, ' $P \rightarrow R$ '), we additionally assume the antecedent to that conditional. Thus our main proof starts:

1	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">$P \rightarrow Q$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">2</td> <td style="padding: 0 5px;">$Q \rightarrow R$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">3</td> <td style="padding: 0 5px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">P</td> </tr> </table> </td> <td></td> </tr> </table>		$P \rightarrow Q$		2	$Q \rightarrow R$		3	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">P</td> </tr> </table>		P	
	$P \rightarrow Q$											
2	$Q \rightarrow R$											
3	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="padding: 0 5px;">P</td> </tr> </table>		P									
	P											

Note that we have made ' P ' available, by treating it as an additional assumption. But now, we can use $\rightarrow E$ on the first premise. This will yield ' Q '. And we can then use $\rightarrow E$ on the second premise. So, by assuming ' P ' we were able to prove ' R ', so we apply the $\rightarrow I$ rule – discharging ' P ' – and finish the proof. Putting all this together, we have:

1	$P \rightarrow Q$	
2	$Q \rightarrow R$	
3	P	
4	Q	$\rightarrow E$ 1, 3
5	R	$\rightarrow E$ 2, 4
6	$P \rightarrow R$	$\rightarrow I$ 3-5

27.4 Additional assumptions and subproofs

The rule $\rightarrow I$ invoked the idea of making additional assumptions. These need to be handled with some care.

Consider this proof:

1	A	
2	B	
3	$B \wedge B$	$\wedge I$ 2, 2
4	B	$\wedge E$ 3
5	$B \rightarrow B$	$\rightarrow I$ 2-4

This is perfectly in keeping with the rules we have laid down already. And it should not seem particularly strange. Since ' $B \rightarrow B$ ' is a tautology, no particular premises should be required to prove it.

But suppose we now tried to continue the proof as follows:

1	A	
2	B	
3	$B \wedge B$	$\wedge I$ 2, 2
4	B	$\wedge E$ 3
5	$B \rightarrow B$	$\rightarrow I$ 2-4
6	B	naughty attempt to invoke $\rightarrow E$ 5, 4

If we were allowed to do this, it would be a disaster. It would allow us to prove any atomic sentence letter from any other atomic sentence letter. But if you tell me that Anne is fast (symbolised by 'A'), I shouldn't be able to conclude that Queen Boudica stood twenty-feet tall (symbolised by 'B')! So we must be prohibited from doing this. But how are we to implement the prohibition?

We can describe the process of making an additional assumption as one of performing a *subproof*: a subsidiary 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 the assumption upon which the subproof will be based. A subproof can be thought of as essentially posing this question: *what could we show, if we also make this additional assumption?*

When we are working within the subproof, we can refer to the additional assumption that we made in introducing the subproof, and to anything that we obtained from our original assumptions. (After all, those original assumptions are still in effect.) But at some point, we shall want to stop working with the additional assumption: we shall want to return from the subproof to the main proof. To indicate that we have returned to the main proof, the vertical line for the subproof comes to an end. At this point, we say that the subproof is CLOSED. Having closed a subproof, we have set aside the additional assumption, so it will be illegitimate to draw upon anything that depends upon that additional assumption. Thus we stipulate:

Any rule whose citation requires mentioning individual lines can mention any earlier lines, *except* for those lines which occur within a closed subproof.

This stipulation rules out the disastrous attempted proof above. The rule of \rightarrow E requires that we cite two individual lines from earlier in the proof. In the purported proof, above, one of these lines (namely, line 4) occurs within a subproof that has (by line 6) been closed. This is illegitimate.

Closing a subproof is called DISCHARGING the assumptions of that subproof. So we can put the point this way: *you cannot refer back to anything that was obtained using discharged assumptions.*

Subproofs, then, allow us to think about what we could show, if we made additional assumptions. The point to take away from this is not surprising – in the course of a proof, we have to keep very careful track of what assumptions we are making, at any given moment. Our proof system does this very graphically. (Indeed, that's precisely why we have chosen to use *this* proof system.)

Once we have started thinking about what we can show by making additional assumptions, nothing stops us from posing the question of what we could show if we were to make *even more* assumptions? This might motivate us to introduce a subproof within a subproof. Here is an example which only uses the rules of proof that we have considered so far:

1	A	
2	B	
3	C	
4	A ∧ B	∧I 1, 2
5	C → (A ∧ B)	→I 3-4
6	B → (C → (A ∧ B))	→I 2-5

Notice that the citation on line 4 refers back to the initial assumption (on line 1) and an assumption of a subproof (on line 2). This is perfectly in order, since neither assumption has been discharged at the time (i.e., by line 4).

Again, though, we need to keep careful track of what we are assuming at any given moment. For suppose we tried to continue the proof as follows:

1	A	
2	B	
3	C	
4	A ∧ B	∧I 1, 2
5	C → (A ∧ B)	→I 3-4
6	B → (C → (A ∧ B))	→I 2-5
7	C → (A ∧ B)	naughty attempt to invoke →I 3-4

This would be awful. If I tell you that Anne is smart, you should not be able to derive that, if Cath is smart (symbolised by 'C') then *both* Anne is smart and Queen Boudica stood 20-feet tall! But this is just what such a proof would suggest, if it were permissible.

The essential problem is that the subproof that began with the assumption 'C' depended crucially on the fact that we had assumed 'B' on line 2. By line 6, we have *discharged* the assumption 'B': we have stopped asking ourselves what we could show, if we also assumed 'B'. So it is simply cheating, to try to help ourselves (on line 7) to the subproof that began with the assumption 'C'. Thus we stipulate, much as before:

Any rule whose citation requires mentioning an entire subproof can mention any earlier subproof, *except* for those subproofs which occur within some *other* closed subproof.

The attempted disastrous proof violates this stipulation. The subproof of lines 3–4 occurs within a subproof that ends on line 5. So it cannot be invoked in line 7.

It is always permissible to open a subproof with any assumption. However, 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 obtain a conditional by \rightarrow I, for instance, you must assume the antecedent of the conditional in a subproof.

Equally, it is always permissible to close a subproof and discharge its assumptions. However, it will not be helpful to do so, until you have reached something useful.

Recall the proof of the argument

$$P \rightarrow Q, Q \rightarrow R \therefore P \rightarrow R$$

from page 162. One thing to note about the proof there is that because there are two assumptions with the same range in the main proof, it is not easily possible to discharge just one of them using the \rightarrow I rule. For that rule only applies to a one-assumption subproof. If we wanted to discharge another of our assumptions, we shall have to put the proof into the right form, with each assumption made individually as the head of its own subproof:

1	$P \rightarrow Q$	
2	$Q \rightarrow R$	
3	P	
4	Q	\rightarrow E 1, 3
5	R	\rightarrow E 2, 4
6	$P \rightarrow R$	\rightarrow I 3–5

The conclusion is now in the range of both assumptions, as in the earlier proof – but now it is also possible to discharge these assumptions if we wish:

1	$P \rightarrow Q$	
2	$Q \rightarrow R$	
3	P	
4	Q	\rightarrow E 1, 3
5	R	\rightarrow E 2, 4
6	$P \rightarrow R$	\rightarrow I 3–5
7	$(Q \rightarrow R) \rightarrow (P \rightarrow R)$	\rightarrow I 2–6
8	$(P \rightarrow Q) \rightarrow ((Q \rightarrow R) \rightarrow (P \rightarrow R))$	\rightarrow I 1–7

While it is permissible, and often convenient, to have several assumptions with the same range and without nesting, I recommend always trying to construct your proofs so that each assumption begins its own subproof. That way, if you later wish to apply rules which discharge a single assumption, you may always do so.

27.5 Conditional introduction and the English conditional

We motivated the conditional introduction rule back on page 160 by giving a English argument, using the English word ‘if’. Now, \rightarrow I is a stipulated rule for our conditional connective \rightarrow ; it doesn’t really need motivation since we could simply postulate that such a rule is part of our formal proof system. It is justified, if justification is needed, by the Deduction Theorem (a result noted in §11.7). (Likewise, \rightarrow E may be justified by a schematic truth table demonstration that $\mathcal{A}, \mathcal{A} \rightarrow \mathcal{C} \vDash \mathcal{C}$.)

But if we are to offer a motivation for our rule in English, then we must be relying on the plausibility of this English analog of \rightarrow I, known as **CONDITIONAL PROOF**:

If you can establish \mathcal{C} , given the assumption that \mathcal{A} and perhaps some supplementary assumptions $\mathcal{B}_1, \dots, \mathcal{B}_n$, then you can establish ‘if \mathcal{A} , \mathcal{C} ’ solely on the basis of the assumptions $\mathcal{B}_1, \dots, \mathcal{B}_n$.

Conditional proof captures a significant aspect of the English ‘if’: the way that conditional helps us neatly summarise reasoning from assumptions, and then store that reasoning for later use in a conditional form.

But if conditional proof is a good rule for English ‘if’, then we can argue that ‘if’ is actually synonymous with ‘ \rightarrow ’:

Suppose that $\mathcal{A} \rightarrow \mathcal{C}$. Assume \mathcal{A} . We can now derive \mathcal{C} , by \rightarrow E. But we have now established \mathcal{C} on the basis of the assumption \mathcal{A} , together with the supplementary assumption $\mathcal{A} \rightarrow \mathcal{C}$. By conditional proof, then, we can establish ‘if \mathcal{A} , \mathcal{C} ’ on the basis of the supplementary assumption $\mathcal{A} \rightarrow \mathcal{C}$ alone. But that of course means that we can derive (in English) an English conditional sentence from a Sentential conditional sentence, with the appropriate interpretation of the constituents \mathcal{A} and \mathcal{C} . Since the English conditional sentence obviously suffices for the Sentential conditional sentence, we have shown them to be synonymous.

Our discussion in §11.5 seemed to indicate that the English conditional was not synonymous with ‘ \rightarrow ’. That is hard to reconcile with the above argument. Most philosophers have concluded that, contrary to appearances, conditional proof is not always a good way of reasoning for English ‘if’. Here is an example which shows this, though it requires some background.

It is clearly valid in English, though rather pointless, to argue from a claim to itself. So when C is some English sentence, $C \therefore C$ is a valid argument in English. And we cannot make a valid argument invalid by adding more premises: if premises we already have are conclusive grounds for the conclusion, adding more premises while keeping the conclusive grounds cannot make the argument less conclusive. So $C, A \therefore C$ is a valid argument in English.

If conditional proof were a good way of arguing in English, we could convert this valid argument into another valid argument with this form: $C \therefore \text{if } A, C$. But then conditional proof would then allow us to convert this valid argument:

1. I will go skiing tomorrow;
 2. I break my leg tonight;
- So: I will go skiing tomorrow. (which is just repeating 1 again)

into this intuitively invalid argument:

1. I will go skiing tomorrow;
- So: If I break my leg tonight, I will go skiing tomorrow.

This is invalid, because even if the premise is true, the conclusion seems to be actually false. If conditional proof enables us to convert a valid English argument into an invalid argument, so much the worse for conditional proof.

There is much more to be said about this example. What is beyond doubt is that \rightarrow I is a good rule for Sentential, regardless of the fortunes of conditional proof in English. It does seem that instances of conditional proof in mathematical reasoning are all acceptable, which again shows the roots of natural deduction as a formalisation of existing mathematical practice. This would suggest that \rightarrow might be a good representation of mathematical uses of ‘if’.

27.6 Biconditional

The rules for the biconditional will be like double-barrelled versions of the rules for the conditional.

In order to prove ‘ $W \leftrightarrow X$ ’, for instance, you must be able to prove ‘ X ’ on the assumption ‘ W ’ and prove ‘ W ’ on the assumption ‘ X ’. The BICONDITIONAL INTRODUCTION rule (\leftrightarrow I) therefore requires two subproofs. Schematically, the rule works like this:

i	\mathcal{A}	
	\vdots	
j	\mathcal{B}	
k	\mathcal{B}	
	\vdots	
l	\mathcal{A}	
	\vdots	
	$\mathcal{A} \leftrightarrow \mathcal{B}$	$\leftrightarrow I\ i-j, k-l$

There can be as many lines as you like between i and j , and as many lines as you like between k and l . Moreover, the subproofs can come in any order, and the second subproof does not need to come immediately after the first. Again, this rule permits us to discharge assumptions, and the same restrictions on making use of claims derived in a closed subproof outside of that subproof apply.

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 obtain the right-hand subsentence. If you have the right-hand subsentence, you can obtain the left-hand subsentence. So we have these two instances of the rule:

m	$\mathcal{A} \leftrightarrow \mathcal{B}$		m	$\mathcal{A} \leftrightarrow \mathcal{B}$	
	\vdots			\vdots	
n	\mathcal{A}		n	\mathcal{B}	
	\vdots			\vdots	
	\mathcal{B}	$\leftrightarrow E\ m, n$		\mathcal{A}	$\leftrightarrow E\ m, n$

Note that the biconditional, and the right or left half, can be separated from one another, and they can appear in any order. However, in the citation for $\leftrightarrow E$, we always cite the biconditional first.

The goodness of our proof rules – that they will never lead us from truth to falsehood, is demonstrated by the following facts, each of which may be demonstrated by consideration of schematic truth tables:

- › $\mathcal{A} \leftrightarrow \mathcal{B}, \mathcal{A} \vDash \mathcal{B}$ and $\mathcal{B} \leftrightarrow \mathcal{A}, \mathcal{B} \vDash \mathcal{A}$;

› If $C_1, \dots, C_n, \mathcal{A} \vDash \mathcal{B}$ and $C_1, \dots, C_n, \mathcal{B} \vDash \mathcal{A}$, then $C_1, \dots, C_n \vDash \mathcal{A} \leftrightarrow \mathcal{B}$.

27.7 Disjunction

Suppose Ludwig is reactionary. Then Ludwig is either reactionary or libertarian. After all, to say that Ludwig is either reactionary or libertarian is to say something weaker than to say that Ludwig is reactionary.

Let me emphasise this point. Suppose Ludwig is reactionary. It follows that Ludwig is *either* reactionary or a kumquat. Equally, it follows that *either* Ludwig is reactionary or that kumquats are the only fruit. Equally, it follows that *either* Ludwig is reactionary or that God is dead. Many of these things are strange inferences to draw. But there is nothing *logically* wrong with them (even if they maybe violate all sorts of implicit conversational norms).

Armed with all this, I present the DISJUNCTION INTRODUCTION rule(s):

$ \begin{array}{l l} m & \mathcal{A} \\ & \vdots \\ & \mathcal{A} \vee \mathcal{B} \quad \vee I m \end{array} $	$ \begin{array}{l l} m & \mathcal{A} \\ & \vdots \\ & \mathcal{B} \vee \mathcal{A} \quad \vee I m \end{array} $
---	---

Notice that \mathcal{B} can be *any* sentence whatsoever. So the following is a perfectly good proof:

1	M	
2	$M \vee \left(((A \leftrightarrow B) \rightarrow (C \wedge D)) \leftrightarrow (E \wedge F) \right)$	$\vee I 1$

Using a truth table to show this would have taken 128 lines.

The disjunction elimination rule is, though, slightly trickier. Suppose that either Ludwig is reactionary or he is libertarian. What can you conclude? Not that Ludwig is reactionary; it might be that he is libertarian instead. And equally, not that Ludwig is libertarian; for he might merely be reactionary. Disjunctions, just by themselves, are hard to work with.

But suppose that we could somehow show both of the following: first, that Ludwig's being reactionary entails that he is an Austrian economist; second, that Ludwig's being libertarian entails that he is an Austrian economist. Then if we know that Ludwig is either reactionary or libertarian, then we know that, whichever he is, Ludwig is an Austrian economist.² This insight can be expressed in the following rule, which is our DISJUNCTION ELIMINATION ($\vee E$) rule:

² This we might call 'no matter whether' reasoning: if each of \mathcal{A} and \mathcal{B} imply \mathcal{C} , then *no matter whether* \mathcal{A} or \mathcal{B} , still \mathcal{C} .

m	$\mathcal{A} \vee \mathcal{B}$					
i	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">⋮</td> <td></td> </tr> </table>	\mathcal{A}		⋮		
\mathcal{A}						
⋮						
j	\mathcal{C}					
k	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{B}</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">⋮</td> <td></td> </tr> </table>	\mathcal{B}		⋮		
\mathcal{B}						
⋮						
l	\mathcal{C}					
	⋮					
	\mathcal{C}	$\vee E\ m, i-j, k-l$				

This is obviously a bit clunkier to write down than our previous rules, but the point is fairly simple. Suppose we have some disjunction, $\mathcal{A} \vee \mathcal{B}$. Suppose we have two subproofs, showing us that \mathcal{C} follows from the assumption that \mathcal{A} , and that \mathcal{C} follows from the assumption that \mathcal{B} . Then we can derive \mathcal{C} itself. As usual, there can be as many lines as you like between i and j , and as many lines as you like between k and l . Moreover, the subproofs and the disjunction can come in any order, and do not have to be adjacent.

Some examples might help illustrate this. Consider this argument:

$$(P \wedge Q) \vee (P \wedge R) \therefore P.$$

An example proof might run thus:

1	$(P \wedge Q) \vee (P \wedge R)$			
2	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$P \wedge Q$</td> <td></td> </tr> </table>	$P \wedge Q$		
$P \wedge Q$				
3	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">P</td> <td style="padding-left: 10px;">$\wedge E\ 2$</td> </tr> </table>	P	$\wedge E\ 2$	
P	$\wedge E\ 2$			
4	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$P \wedge R$</td> <td></td> </tr> </table>	$P \wedge R$		
$P \wedge R$				
5	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">P</td> <td style="padding-left: 10px;">$\wedge E\ 4$</td> </tr> </table>	P	$\wedge E\ 4$	
P	$\wedge E\ 4$			
6	P	$\vee E\ 1, 2-3, 4-5$		

Here is a slightly harder example. Consider:

$$A \wedge (B \vee C) \therefore (A \wedge B) \vee (A \wedge C).$$

Here is a proof corresponding to this argument:

1	$A \wedge (B \vee C)$	
2	A	$\wedge E$ 1
3	$B \vee C$	$\wedge E$ 1
4	B	
5	$A \wedge B$	$\wedge I$ 2, 4
6	$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 5
7	C	
8	$A \wedge C$	$\wedge I$ 2, 7
9	$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 8
10	$(A \wedge B) \vee (A \wedge C)$	$\vee E$ 3, 4–6, 7–9

Don't be alarmed if you think that you wouldn't have been able to come up with this proof yourself. The ability to come up with novel proofs will come with practice. The key question at this stage is whether, looking at the proof, you can see that it conforms with the rules that we have laid down. And that just involves checking every line, and making sure that it is justified in accordance with the rules we have laid down.

Again, our disjunction rules are supported by some valid Sentential argument forms:

- › $\mathcal{A} \models \mathcal{A} \vee B$ and $B \models \mathcal{A} \vee B$;
- › If $\mathcal{D}_1, \dots, \mathcal{D}_n, \mathcal{A} \vee B, \mathcal{A} \models C$ and $\mathcal{D}_1, \dots, \mathcal{D}_n, \mathcal{A} \vee B, B \models C$, then $\mathcal{D}_1, \dots, \mathcal{D}_n, \mathcal{A} \vee B \models C$.

27.8 Negation

Natural deduction rules for negation are tricky, in part because there are so many options. All of them stem from the same basic insight: when assuming \mathcal{A} goes awry in some way, conclude $\neg \mathcal{A}$; and likewise, if assuming $\neg \mathcal{A}$ goes awry, conclude \mathcal{A} .

Our NEGATION INTRODUCTION rule fits this pattern very clearly:

i	\mathcal{A}	
	\vdots	
j	\mathcal{B}	
	\vdots	
k	$\neg\mathcal{B}$	
	\vdots	
	$\neg\mathcal{A}$	$\neg I\ i-j, i-k$

Here, we can derive a sentence and its negation both within the range of the assumption that \mathcal{A} . So if \mathcal{A} were assumed, something contradictory would be derivable under that assumption. (We could apply conjunction introduction to lines j and k to make the contradiction explicit, but that wouldn't be strictly necessary.) Since contradictions are fundamentally unacceptable as the termination of a chain of argument, we must have begun with an inappropriate starting point when we assumed \mathcal{A} . So, in fact, we conclude, $\neg\mathcal{A}$, discharging our erroneous assumption that \mathcal{A} . There is no need for the line with \mathcal{B} on it to occur before the line with $\neg\mathcal{B}$ on it.

This rule is interesting, because it is *almost* its own elimination rule too! Consider this schematic proof:

i	$\neg\mathcal{A}$	
	\vdots	
j	\mathcal{B}	
	\vdots	
k	$\neg\mathcal{B}$	
$k + 1$	$\neg\neg\mathcal{A}$	$\neg I\ i-j, i-k$

This proof terminates with a sentence that is logically equivalent to \mathcal{A} , discharging the assumption that $\neg\mathcal{A}$ because it leads to contradictory conclusions. This looks awfully close to a rule of negation elimination – if only we could find a way to replace a doubly-negated sentence $\neg\neg\mathcal{A}$ by the logically equivalent sentence \mathcal{A} , which would have eliminated the negation from the problematic assumption $\neg\mathcal{A}$.

In our system, we approach this problem by the brute force method. This is our rule of DOUBLE NEGATION ELIMINATION:

i	$\neg\neg\mathcal{A}$	
	\vdots	
	\mathcal{A}	$\neg\neg E\ i$

In Sentential, the double negation $\neg\neg\mathcal{A}$ is equivalent to \mathcal{A} . In natural languages, too, double negations tend to cancel out. That said, you should be aware that context and emphasis can prevent them from doing so. Consider: ‘Jane is not *not* happy’. Arguably, one cannot derive ‘Jane is happy’, since the first sentence should be understood as meaning the same as ‘Jane is not *unhappy*’. This is compatible with ‘Jane is in a state of profound indifference’. As usual, moving to Sentential forces us to sacrifice certain nuances of English expressions – we have, in Sentential, just one resource for translating negative expressions like ‘not’ and the suffix ‘un-’, even if they are not synonyms in English.

If we apply $\neg\neg E$ to the last line of the preceding schematic proof, we get a proof that shows we can eliminate a negation from a negated sentence, so long as that sentence leads to trouble of the appropriate sort:

i	$\neg\mathcal{A}$	
	\vdots	
j	\mathcal{B}	
	\vdots	
k	$\neg\mathcal{B}$	
$k + 1$	$\neg\neg\mathcal{A}$	$\neg I\ i-j, i-k$
$k + 2$	\mathcal{A}	$\neg\neg E\ k + 1$

This proof shows that, in a system with $\neg I$ and $\neg\neg E$, we do not need any separate elimination rule for a single negation – the effect of any such rule could be perfectly simulated by the above schematic proof. The addition of a single negation elimination rule would not allow us to prove any more than we already can. We will see some other rules with this conservative impact later in this chapter.

To see our negation rules in action, consider:

$$P \therefore (P \wedge D) \vee (P \wedge \neg D).$$

Here is a proof corresponding with the argument:

1	P							
2	$\neg((P \wedge D) \vee (P \wedge \neg D))$							
3	<table style="border-collapse: collapse; margin-left: 5px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px; border-bottom: 1px solid black;">D</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$P \wedge D$</td> <td style="padding-left: 10px;">$\rightarrow I$ 1, 3</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$(P \wedge D) \vee (P \wedge \neg D)$</td> <td style="padding-left: 10px;">$\vee I$ 4</td> </tr> </table>	D		$P \wedge D$	$\rightarrow I$ 1, 3	$(P \wedge D) \vee (P \wedge \neg D)$	$\vee I$ 4	
D								
$P \wedge D$	$\rightarrow I$ 1, 3							
$(P \wedge D) \vee (P \wedge \neg D)$	$\vee I$ 4							
4	$P \wedge D$	$\rightarrow I$ 1, 3						
5	$(P \wedge D) \vee (P \wedge \neg D)$	$\vee I$ 4						
6	$\neg D$	$\neg I$ 3-5, 3-2						
7	$P \wedge \neg D$	$\rightarrow I$ 1, 6						
8	$(P \wedge D) \vee (P \wedge \neg D)$	$\vee I$ 7						
9	$\neg\neg((P \wedge D) \vee (P \wedge \neg D))$	$\neg I$ 2-8						
10	$(P \wedge D) \vee (P \wedge \neg D)$	$\neg\neg E$ 9						

I make two comments. In line 6, the justification cites line 2 which lies outside the subproof. That is okay, since the application of the rule lies within the range of the assumption of line 2. In line 9, the justification only cites the subproof from 2 to 8, rather than two ranges of line numbers. This is because in this application of our rule, we have the special case where the sentence such that both it and its negation can be derived from the assumption is that assumption. It would be trivial to derive it from itself.

The negation rules are supported by these valid Sentential argument forms:

- › $\vee \vee \mathcal{A} \vDash \mathcal{A}$;
- › If $\mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{A} \vDash \mathcal{B}$ and $\mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{A} \vDash \neg\mathcal{B}$, then $\mathcal{C}_1, \dots, \mathcal{C}_n \vDash \neg\mathcal{A}$.

These are all of the basic rules for the proof system for Sentential. Figure 27.1 shows a long proof involving most of our rules in action.

27.9 Inferentialism

You will have noticed that our rules come in pairs: an introduction rule that tells you how to introduce a connective into a proof from what you have already, and an elimination rule that tells you how to remove it in favour of its consequences. These rules can be justified by consideration of the meanings we assigned to the connectives of Sentential in the schematic truth tables of §8.

But perhaps we should invert this order of justification. After all, the proof rules already seem to capture (more or less) how ‘and’, ‘or’, ‘not’, ‘if’, and ‘iff’ work in conversation. We make some claims and assumptions. The introduction and elimination rules summarise

1	$\neg P \vee Q$	
2	$\neg P$	
3	P	
4	$\neg Q$	
5	$P \wedge \neg P$	$\wedge I$ 3, 2
6	P	$\wedge E$ 5
7	$\neg P$	$\wedge E$ 5
8	$\neg\neg Q$	$\neg I$ 4-7
9	Q	$\neg E$ 8
10	$P \rightarrow Q$	$\rightarrow I$ 3-9
11	Q	
12	P	
13	$Q \wedge Q$	$\wedge I$ 11, 11
14	Q	$\wedge E$ 13
15	$P \rightarrow Q$	$\rightarrow I$ 12-14
16	$P \rightarrow Q$	$\vee E$ 1, 2-10, 11-15
17	$P \rightarrow Q$	
18	$\neg(\neg P \vee Q)$	
19	P	
20	Q	$\rightarrow E$ 17, 19
21	$\neg P \vee Q$	$\vee I$ 20
22	$\neg P$	$\neg I$ 19-18, 19-21
23	$\neg P \vee Q$	$\vee I$ 22
24	$\neg\neg\neg P \vee Q$	$\neg I$ 18-18, 18-23
25	$\neg P \vee Q$	$\neg E$ 24
26	$(\neg P \vee Q) \leftrightarrow (P \rightarrow Q)$	$\leftrightarrow I$ 1-16, 17-25

Figure 27.1: Proof from no assumptions of $(\neg P \vee Q) \leftrightarrow (P \rightarrow Q)$

how we might proceed in our conversation against the background of those claims and assumptions. Many philosophers have thought that the meaning of an expression is entirely governed by how it is or might be used in a conversation by competent speakers – in slogan form, *meaning is fixed by use*. If these proof rules describe how we might bring an expression into a conversation, and what we may do with it once it is there, then these proof rules describe the totality of facts on which meaning depends. The meaning of a connective, according to this INFERENTIALIST picture, is represented by its introduction and elimination rules – and not by the truth-function that a schematic truth table represents. On this view, it is the correctness of the schematic proof of $\mathcal{A} \vee \mathcal{B}$ from \mathcal{A} which explains *why* the schematic truth table for ‘ \vee ’ has a T on every row on which at least one of its constituents gets a T.

There is a significant debate on just this issue in the philosophy of language, about the nature of *meaning*. Is the meaning of a word what it represents, the view sometimes called REPRESENTATIONALISM? Or is the meaning of a word, rather, given by some rules for how to use it, as inferentialism says? We cannot go deeply into this issue here, but I will say a little. The representationalist view seems to accomodate some expressions very well: the meaning of a name, for example, seems very plausibly to be identified with what it names; the meaning of a predicate might be thought of as the corresponding property. But inferentialism seems more natural as an approach to the logical connectives:

Anyone who has learnt to perform [conjunction introduction and conjunction elimination] knows the meaning of ‘and’, for there is simply nothing more to knowing the meaning of ‘and’ than being able to perform these inferences.³

It seems rather unnatural, by contrast, to think that the meaning of ‘and’ is some abstract mathematical ‘thing’ represented by a truth table.

Can the inferentialist distinguish good systems of rules, such as those governing ‘and’, from bad systems? The problem is that without appealing to truth tables or the like, we seem to be committed to the legitimacy of rather problematic connectives. The most famous example is Prior’s ‘tonk’ governed by these rules:

m	\mathcal{A}		m	$\mathcal{A} \text{ tonk } \mathcal{B}$		n	\mathcal{B}
	\vdots			\vdots			
n	$\mathcal{A} \text{ tonk } \mathcal{B}$	tonk-I m	n	\mathcal{B}	tonk-E m		

You will notice that ‘tonk’ has an introduction rule like ‘ \vee ’, and an elimination rule like ‘ \wedge ’. Of course ‘tonk’ is a connective we would not like in a language, since pairing the introduction and elimination rules would allow us to prove any arbitrary sentence from any assumption whatsoever:

³ A N Prior, ‘The Runabout Inference-Ticket’, *Analysis* 21 (1961), p. 38.

1	P	
2	$P \text{ tonk } Q$	tonk-I 1
3	Q	tonk-E 2

If we are to rule out such deviant connectives as ‘tonk’, Prior argues, we have to accept that ‘an expression must have some independently determined meaning before we can discover whether inferences involving it are valid or invalid’ (Prior, *op. cit.*, p. 38). We cannot, that is, accept the inferentialist position that the rules of implication come first and the meaning comes second. Inferentialists have replied, but we must unfortunately leave this interesting debate here for now.⁴

Practice exercises

A. The following two ‘proofs’ are *incorrect*. Explain the mistakes they make.

1	$\neg L \rightarrow (A \wedge L)$		1	$A \wedge (B \wedge C)$	
2	$\neg L$		2	$(B \vee C) \rightarrow D$	
3	A	$\rightarrow E$ 1, 2	3	B	$\wedge E$ 1
4	L		4	$B \vee C$	$\vee I$ 3
5	$L \wedge \neg L$	$\perp I$ 4, 2	5	D	$\rightarrow E$ 4, 2
6	$\neg A$				
7	L	$\rightarrow I$ 5			
8	$\neg L$	$\rightarrow E$ 5			
9	A	$\neg\neg E$ 6-7, 6-8			
10	A	$\vee E$ 2-3, 4-9			

B. The following three proofs are missing their citations (rule and line numbers). Add them, to turn them into bona fide proofs. Additionally, write down the argument that corresponds to each proof.

⁴ The interested reader might wish to start with this reply to Prior: Nuel D Belnap, Jr, ‘Tonk, Plonk and Plink’, *Analysis* 22 (1962), pp. 130-34.

1	$P \wedge S$
2	P
3	S
4	$S \rightarrow R$
5	R
6	$R \vee E$

1	$A \rightarrow D$
2	$A \wedge B$
3	A
4	D
5	$D \vee E$
6	$(A \wedge B) \rightarrow (D \vee E)$

1	$\neg L \rightarrow (J \vee L)$
2	$\neg L$
3	$J \vee L$
4	J
5	$J \wedge J$
6	J
7	L
8	$\neg J$
9	$\neg \neg J$
10	J
11	J

C. Give a proof for each of the following arguments:

1. $J \rightarrow \neg J \therefore \neg J$
2. $Q \rightarrow (Q \wedge \neg Q) \therefore \neg Q$
3. $A \rightarrow (B \rightarrow C) \therefore (A \wedge B) \rightarrow C$
4. $K \wedge L \therefore K \leftrightarrow L$
5. $(C \wedge D) \vee E \therefore E \vee D$
6. $A \leftrightarrow B, B \leftrightarrow C \therefore A \leftrightarrow C$
7. $\neg F \rightarrow G, F \rightarrow H \therefore G \vee H$
8. $(Z \wedge K) \vee (K \wedge M), K \rightarrow D \therefore D$
9. $P \wedge (Q \vee R), P \rightarrow \neg R \therefore Q \vee E$
10. $S \leftrightarrow T \therefore S \leftrightarrow (T \vee S)$
11. $\neg(P \rightarrow Q) \therefore \neg Q$
12. $\neg(P \rightarrow Q) \therefore P$

28 Proof-theoretic concepts

We shall introduce some new vocabulary. The following expression:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash \mathcal{C}$$

means that there is some proof which starts with assumptions among $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ and ends with \mathcal{C} (and no undischarged assumptions other than those we started with). Derivatively, we shall write:

$$\vdash \mathcal{A}$$

to mean that there is a proof of \mathcal{A} which ends up having no undischarged assumptions.

The symbol ‘ \vdash ’ is called the *single turnstile*. I want to emphasise that this is not the double turnstile symbol (‘ \vDash ’) that we used to symbolise entailment in chapters 3 and 5. The single turnstile, ‘ \vdash ’, concerns the existence of proofs; the double turnstile, ‘ \vDash ’, concerns the existence of valuations (or interpretations, when used for Quantifier). *They are very different notions.*¹

Armed with our ‘ \vdash ’ symbol, we can introduce some new terminology.

\mathcal{A} is a THEOREM iff $\vdash \mathcal{A}$

To illustrate this, suppose I want to prove that ‘ $\neg(A \wedge \neg A)$ ’ is a theorem. So I must start my proof without *any* assumptions. However, since I want to prove a sentence whose main connective is a negation, I shall want to immediately begin a subproof, with the additional assumption ‘ $A \wedge \neg A$ ’, and show that this leads to contradictory consequences. All told, then, the proof looks like this:

1	$A \wedge \neg A$	
2	A	$\wedge E$ 1
3	$\neg A$	$\wedge E$ 1
4	$\neg(A \wedge \neg A)$	$\neg I$ 1-2, 1-3

¹ They are very different, but if we’ve designed our proof system well, we shouldn’t be able to prove a conclusion from some assumptions *unless* that conclusion validly follows from those assumptions. And if we are really fortunate, we should be able to provide a proof corresponding to any valid argument. More on this in §34.

We have therefore proved ‘ $\neg(A \wedge \neg A)$ ’ on no (undischarged) assumptions. This particular theorem is an instance of what is sometimes called *the Law of Non-Contradiction*, that for any \mathcal{A} , $\neg(\mathcal{A} \wedge \neg\mathcal{A})$.

To show that something is a theorem, you just have to find a suitable proof. It is typically much harder to show that something is *not* a theorem. To do 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. Perhaps you just didn’t try hard enough.

Here is another new bit of terminology:

Two sentences \mathcal{A} and \mathcal{B} are PROVABLY EQUIVALENT iff each can be proved from the other; i.e., both $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$.

As in the case of showing that a sentence is a theorem, 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 is just as hard as showing that a sentence is not a theorem. Notice that if we have two proofs, one demonstrating $\mathcal{A} \vdash \mathcal{B}$ and the other demonstrating $\mathcal{B} \vdash \mathcal{A}$, we can put them together as subproofs of a single proof, and apply \leftrightarrow I to provide a proof from no assumptions which demonstrates that $\vdash \mathcal{A} \leftrightarrow \mathcal{B}$.

Here is a third, related, bit of terminology:

The sentences $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are JOINTLY CONTRARY iff a contradiction can be proved from them, i.e., for some \mathcal{B} , $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash \mathcal{B} \wedge \neg\mathcal{B}$.

It is easy to show that some sentences are jointly contrary: you just need to provide a proof of the contradiction from assuming all the sentences. Showing that some sentences are not jointly contrary is much harder. It would require more than just providing a proof or two; it would require showing that no proof of a certain kind is *possible*.

This table summarises whether one or two proofs suffice, or whether we must reason about all possible proofs.

	Yes	No
theorem?	one proof	all possible proofs
equivalent?	two proofs	all possible proofs
not contrary?	all possible proofs	one proof

Practice exercises

A. Show that each of the following sentences is a theorem:

1. $O \rightarrow O$
2. $N \vee \neg N$
3. $J \leftrightarrow (J \vee (L \wedge \neg L))$
4. $((A \rightarrow B) \rightarrow A) \rightarrow A$

B. Provide proofs to show each of the following:

1. $C \rightarrow (E \wedge G), \neg C \rightarrow G \vdash G$
2. $M \wedge (\neg N \rightarrow \neg M) \vdash (N \wedge M) \vee \neg M$
3. $(Z \wedge K) \leftrightarrow (Y \wedge M), D \wedge (D \rightarrow M) \vdash Y \rightarrow Z$
4. $(W \vee X) \vee (Y \vee Z), X \rightarrow Y, \neg Z \vdash W \vee Y$

C. Show that each of the following pairs of sentences are provably equivalent:

1. $R \leftrightarrow E, E \leftrightarrow R$
2. $G, \neg\neg\neg\neg G$
3. $T \rightarrow S, \neg S \rightarrow \neg T$
4. $U \rightarrow I, \neg(U \wedge \neg I)$
5. $\neg(C \rightarrow D), C \wedge \neg D$
6. $\neg G \leftrightarrow H, \neg(G \leftrightarrow H)$

D. If you know that $\mathcal{A} \vdash \mathcal{B}$, what can you say about $(\mathcal{A} \wedge \mathcal{C}) \vdash \mathcal{B}$? What about $(\mathcal{A} \vee \mathcal{C}) \vdash \mathcal{B}$? Explain your answers.

E. In this section, I claimed that it is just as hard to show that two sentences are not provably equivalent, as it is to show that a sentence is not a theorem. Why did I claim this? (*Hint*: think of a sentence that would be a theorem iff \mathcal{A} and \mathcal{B} were provably equivalent.)

29 Proof strategies

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 obtain the conclusion. Look at the conclusion and ask what the introduction rule is for its main connective. 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 get to 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} . The subproof ought to end with \mathcal{B} . So, what can you do to get \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 obtained so far. Think about the elimination rules for the main operators of these sentences. These will tell you what your options are.

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.

Try proceeding indirectly. If you cannot find a way to show \mathcal{A} directly, try starting by assuming $\neg\mathcal{A}$. If a contradiction follows, then you will be able to obtain $\neg\neg\mathcal{A}$ by \neg I, and then \mathcal{A} by $\neg\neg$ E. This will often be a good way of proceeding when the conclusion you are aiming at has a disjunction as its main connective.

Persist. These are guidelines, not laws. Try different things. If one approach fails, then try something else.

For example: suppose you tried to follow the idea ‘work backwards from what you want’ in establishing ‘ $P, P \rightarrow (P \rightarrow Q) \vdash P \rightarrow Q$ ’. You would be tempted to start a subproof from the assumption ‘ P ’, and while that proof strategy would eventually succeed, you would have done better to simply apply \rightarrow E and terminate after one proof step.

By contrast, suppose you tried to follow the idea ‘work forwards from what you have’ in trying to establish ‘ $(P \vee (Q \vee (R \vee S))) \vdash P \rightarrow P$ ’. You might begin an awkward nested series

of subproofs to apply $\vee E$ to the disjunctive premise. But beginning with the conclusion might prompt you instead to simply open a subproof from the assumption P , and the subsequent proof will make no use of the premise at all, as the conclusion is a theorem.

Neither of these heuristics is sacrosanct. You will get a sense of how to construct proofs efficiently and fluently with practice. Unfortunately there is no quick substitute for practice.

30 Derived rules for Sentential

In §27, we introduced the basic rules of our proof system for Sentential. In this section, we shall consider some alternative or additional rules for our system.

None of these rules adds anything fundamentally to our system. They are all DERIVED rules, which means that anything we can prove by using them, we could have proved using just the rules in our original official system of natural deduction proofs. Any of these rules is a *conservative* addition to our proof system, because none of them would enable us to prove anything we could not already prove. (Adding the rules for ‘tonk’ from §27.9, by contrast, would allow us to prove many new things – any system which includes those rules is not a conservative extension of our original system of proof rules.)

But sometimes adding new rules can shorten proofs, or make them more readable and user-friendly. And some of them are of interest in their own right, as arguably independently plausible rules of implication as they stand, or as alternative rules we could have taken as basic instead.

30.1 Reiteration

The first additional rule is *reiteration* (R). This just allows us to repeat ourselves:

m		\mathcal{A}	
		\mathcal{A}	R m

Such a rule is obviously legitimate; but one might well wonder how such a rule could ever be useful. Well, consider:

1	$\neg(P \vee Q)$	
2	P	
3	$P \vee Q$	$\vee I$ 2
4	$\neg(P \vee Q)$	R 1
5	$\neg P$	$\neg I$ 2-4

This is a fairly typical use of the R rule. If we had not used it, we would have had an awkward (though by no means incorrect) appeal to the instance of ' $\neg(P \vee Q)$ ' in line 1 as part of the application of the $\neg I$ rule to line 2. The use of reiteration makes it obvious that a contradiction really is derivable in the range of the assumption ' P '. The primary use of reiteration is simply to make clear when a sentence falls in the range of an assumption.

But we need not have used a rule of reiteration. We could replace each application of the reiteration rule on some line $k + 1$ (reiterating some prior line m) with the following combination of moves deploying just the *basic* rules of §27:

m	\mathcal{A}	
k	$\mathcal{A} \wedge \mathcal{A}$	$\wedge I$ m
$k + 1$	\mathcal{A}	$\wedge E$ k

To be clear: this is not a proof. Rather, it is a proof *scheme*. After all, it uses a variable, ' \mathcal{A} ', rather than a sentence of Sentential. But the point is simple. Whatever sentences of Sentential we plugged in for ' \mathcal{A} ', and whatever lines we were working on, we could produce a legitimate proof. So you can think of this as a recipe for producing proofs. Indeed, it is a recipe which shows us that, anything we can prove using the rule R, we can prove (with one more line) using just the *basic* rules of §27. So we can describe the rule R as a derived rule, since its justification is derived from our basic rules.

You might note that in lines 5-7 in the complicated proof in Figure 27.1, we in effect made use of this proof scheme, introducing a conjunction from prior lines only to immediately eliminate again, just to ensure that the relevant sentences appeared directly in the range of the assumption ' Q '.

30.2 Disjunctive syllogism

Here is a very natural argument form.

Mitt is either in Massachusetts or in DC. He is not in DC. So, he is in Massachusetts.

This inference pattern is called *disjunctive syllogism*. We could add it to our proof system:

$\begin{array}{l l} m & \mathcal{A} \vee \mathcal{B} \\ n & \neg \mathcal{A} \\ & \mathcal{B} \end{array} \quad \text{DS } m, n$	$\begin{array}{l l} m & \mathcal{A} \vee \mathcal{B} \\ n & \neg \mathcal{B} \\ & \mathcal{A} \end{array} \quad \text{DS } m, n$
--	--

This is, if you like, a new rule of disjunction elimination. But there is nothing fundamentally new here. We can emulate the rule of disjunctive syllogism using our basic proof rules, as this schematic proof indicates:

m	$\mathcal{A} \vee \mathcal{B}$											
n	$\neg \mathcal{A}$											
	\vdots											
k	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{B}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td style="padding-left: 10px;">R k</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{A}$</td> <td style="padding-left: 10px;">R n</td> </tr> </table> </td> <td style="padding-left: 10px;">$\neg \text{I } k+1-k+3$</td> </tr> </table>	\mathcal{A}		<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{B}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td style="padding-left: 10px;">R k</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{A}$</td> <td style="padding-left: 10px;">R n</td> </tr> </table>	$\neg \mathcal{B}$		\mathcal{A}	R k	$\neg \mathcal{A}$	R n	$\neg \text{I } k+1-k+3$	
\mathcal{A}												
<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{B}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">\mathcal{A}</td> <td style="padding-left: 10px;">R k</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg \mathcal{A}$</td> <td style="padding-left: 10px;">R n</td> </tr> </table>	$\neg \mathcal{B}$		\mathcal{A}	R k	$\neg \mathcal{A}$	R n	$\neg \text{I } k+1-k+3$					
$\neg \mathcal{B}$												
\mathcal{A}	R k											
$\neg \mathcal{A}$	R n											
$k+1$												
$k+2$												
$k+3$												
$k+4$	$\neg \mathcal{B}$											
$k+5$	\mathcal{B}	$\neg \text{E } k+4$										
$k+6$	\mathcal{B}											
$k+7$	\mathcal{B}	R $k+6$										
$k+8$	\mathcal{B}	$\vee \text{E } m, k-k+5, k+6-k+7$										

We have used the rule of reiteration in this schematic proof, but we already know that any uses of that rule can themselves be replaced by more roundabout proofs using conjunction introduction and elimination, if required. So adding disjunctive syllogism would not make any new proofs possible that were not already obtainable in our original system.

30.3 Modus tollens

Another useful pattern of inference is embodied in the following argument:

If Hillary won the election, then she is in the White House. She is not in the White House. So she did not win the election.

This inference pattern is called *modus tollens*. The corresponding rule is:

m	$\mathcal{A} \rightarrow \mathcal{B}$	
n	$\neg \mathcal{B}$	
	$\neg \mathcal{A}$	$\text{MT } m, n$

This is, if you like, a new rule of conditional elimination.

This rule is, again, a conservative addition to our stock of proof rules. Any application of it could be emulated by this form of proof using our original rules:

m	$\mathcal{A} \rightarrow \mathcal{B}$	
n	$\neg \mathcal{B}$	
	\vdots	
k	\mathcal{A}	
$k + 1$	\mathcal{B}	$\rightarrow \text{E } m, k$
$k + 2$	$\neg \mathcal{B}$	$\text{R } n$
$k + 3$	$\neg \mathcal{A}$	$\neg \text{I } k-k + 2$

Again, we made a dispensible use of reiteration.

30.4 *Tertium non datur*

Suppose that we can show that if it's sunny outside, then Bill will have brought an umbrella (for fear of burning). Suppose we can also show that, if it's not sunny outside, then Bill will have brought an umbrella (for fear of rain). Well, there is no third way for the weather to be. So, *whatever the weather*, Bill will have brought an umbrella.

This line of thinking motivates the following rule:

i	\mathcal{A}	
j	\mathcal{B}	
k	$\neg\mathcal{A}$	
l	\mathcal{B}	
	\mathcal{B}	TND $i-j, k-l$

The rule is sometimes called *tertium non datur*, which means ‘no third way’. There can be as many lines as you like between i and j , and as many lines as you like between k and l . Moreover, the subproofs can come in any order, and the second subproof does not need to come immediately after the first.

Tertium non datur is able to be emulated using just our original proof rules. Here is a schematic proof which demonstrates this:

i	\mathcal{A}	
j	\mathcal{B}	
k	$\neg\mathcal{A}$	
l	\mathcal{B}	
	\vdots	
m	$\mathcal{A} \rightarrow \mathcal{B}$	$\rightarrow I\ i-j$
$m+1$	$\neg\mathcal{A} \rightarrow \mathcal{B}$	$\rightarrow I\ k-l$
$m+2$	$\neg\mathcal{B}$	
$m+3$	\mathcal{A}	
$m+4$	\mathcal{B}	$\rightarrow E\ m, m+3$
$m+5$	$\neg\mathcal{B}$	R $m+2$
$m+6$	$\neg\mathcal{A}$	$\neg I\ m+3-m+5$
$m+7$	\mathcal{B}	$\rightarrow E\ m+1, m+6$
$m+8$	$\neg\neg\mathcal{B}$	$\neg I\ m+2-m+7$
$m+9$	\mathcal{B}	$\neg E\ m+8$

Again a dispensible use of reiteration occurs in this proof just to make it more readable.

30.5 De Morgan Rules

Our final additional rules are called De Morgan's Laws. (These are named after the nineteenth century logician August De Morgan.) The first two De Morgan rules show the provable equivalence of a negated conjunction and a disjunction of negations.

$m \quad \quad \neg(\mathcal{A} \wedge \mathcal{B})$ $\quad \quad \quad \neg\mathcal{A} \vee \neg\mathcal{B} \quad \text{DeM } m$	$m \quad \quad \neg\mathcal{A} \vee \neg\mathcal{B}$ $\quad \quad \quad \neg(\mathcal{A} \wedge \mathcal{B}) \quad \text{DeM } m$
--	--

The second pair of De Morgan rules are dual to the first pair: they show the provable equivalence of a negated disjunction and a conjunction of negations.

$m \quad \quad \neg(\mathcal{A} \vee \mathcal{B})$ $\quad \quad \quad \neg\mathcal{A} \wedge \neg\mathcal{B} \quad \text{DeM } m$	$m \quad \quad \neg\mathcal{A} \wedge \neg\mathcal{B}$ $\quad \quad \quad \neg(\mathcal{A} \vee \mathcal{B}) \quad \text{DeM } m$
--	--

These are all of the additional rules of our proof system for Sentential.

The De Morgan rules are no genuine addition to the power of our original natural deduction system. Here is a demonstration of how we could derive the first De Morgan rule:

k	$\neg(\mathcal{A} \wedge \mathcal{B})$	
m	$\neg(\neg\mathcal{A} \vee \neg\mathcal{B})$	
$m + 1$	$\quad \quad \neg\mathcal{A}$	
$m + 2$	$\quad \quad \neg\mathcal{A} \vee \neg\mathcal{B}$	$\vee\text{I } m + 1$
$m + 3$	$\neg\neg\mathcal{A}$	$\neg\text{I } m + 1 - m + 2, m + 1 - m$
$m + 4$	\mathcal{A}	$\neg\text{E } m + 3$
$m + 5$	$\quad \quad \neg\mathcal{B}$	
$m + 6$	$\quad \quad \neg\mathcal{A} \vee \neg\mathcal{B}$	$\vee\text{I } m + 5$
$m + 7$	$\neg\neg\mathcal{B}$	$\neg\text{I } m + 5 - m + 6, m + 5 - m$
$m + 8$	\mathcal{B}	$\neg\text{E } m + 7$
$m + 9$	$\mathcal{A} \wedge \mathcal{B}$	$\wedge\text{I } m + 4, m + 8$
$m + 10$	$\neg\neg(\neg\mathcal{A} \vee \neg\mathcal{B})$	$\neg\text{I } m - m + 9, m - k$
$m + 11$	$\neg\mathcal{A} \vee \neg\mathcal{B}$	$\neg\text{E } m + 10$

Here is a demonstration of how we could derive the second De Morgan rule:

k	$\neg\mathcal{A} \vee \neg\mathcal{B}$					
m	$\neg\mathcal{A}$					
$m+1$	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">$\mathcal{A} \wedge \mathcal{B}$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">\mathcal{A}</td> <td></td> </tr> </table>	$\mathcal{A} \wedge \mathcal{B}$		\mathcal{A}		
$\mathcal{A} \wedge \mathcal{B}$						
\mathcal{A}						
$m+2$	\mathcal{A}	$\wedge E\ m+1$				
$m+3$	$\neg(\mathcal{A} \wedge \mathcal{B})$	$\neg I\ m+1-m+2, m+1-m$				
$m+4$	$\neg\mathcal{B}$					
$m+5$	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">$\mathcal{A} \wedge \mathcal{B}$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">\mathcal{B}</td> <td></td> </tr> </table>	$\mathcal{A} \wedge \mathcal{B}$		\mathcal{B}		
$\mathcal{A} \wedge \mathcal{B}$						
\mathcal{B}						
$m+6$	\mathcal{B}	$\wedge E\ m+5$				
$m+7$	$\neg(\mathcal{A} \wedge \mathcal{B})$	$\neg I\ m+5-m+6, m+5-m+4$				
$m+8$	$\neg(\mathcal{A} \wedge \mathcal{B})$	$\vee E\ k, m-m+3, m+4-m+7$				

Similar demonstrations can be offered explaining how we could derive the third and fourth De Morgan rules. These are left as exercises.

Practice exercises

A. The following proofs are missing their citations (rule and line numbers). Add them wherever they are required: you may use any of the original or derived rules, as appropriate.

1	$W \rightarrow \neg B$	1	$L \leftrightarrow \neg O$				
2	$A \wedge W$	2	$L \vee \neg O$				
3	$B \vee (J \wedge K)$	3	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">$\neg L$</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding-right: 5px;">$\neg O$</td> <td></td> </tr> </table>	$\neg L$		$\neg O$	
$\neg L$							
$\neg O$							
4	W	4	L				
5	$\neg B$	5	$\neg L$				
6	$J \wedge K$	6	$\neg\neg L$				
7	K	7	L				
		8	L				

1	$Z \rightarrow (C \wedge \neg N)$	
2	$\neg Z \rightarrow (N \wedge \neg C)$	
3	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(N \vee C)$</td> </tr> </table>	$\neg(N \vee C)$
$\neg(N \vee C)$		
4	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg N \wedge \neg C$</td> </tr> </table>	$\neg N \wedge \neg C$
$\neg N \wedge \neg C$		
5	$\neg N$	
6	$\neg C$	
7	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">Z</td> </tr> </table>	Z
Z		
8	<table style="border-collapse: collapse; margin-left: 20px;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$C \wedge \neg N$</td> </tr> </table>	$C \wedge \neg N$
$C \wedge \neg N$		
9	C	
10	$\neg C$	
11	$\neg Z$	
12	$N \wedge \neg C$	
13	N	
14	$\neg\neg(N \vee C)$	
15	$N \vee C$	

B. Give a proof for each of these arguments:

1. $E \vee F, F \vee G, \neg F \therefore E \wedge G$
2. $M \vee (N \rightarrow M) \therefore \neg M \rightarrow \neg N$
3. $(M \vee N) \wedge (O \vee P), N \rightarrow P, \neg P \therefore M \wedge O$
4. $(X \wedge Y) \vee (X \wedge Z), \neg(X \wedge D), D \vee M \therefore M$

C. Provide proof schemes that justify the addition of the third and fourth De Morgan rules as derived rules.

D. The proofs you offered in response to question A above used derived rules. Replace the use of derived rules, in such proofs, with only basic rules. You will find some ‘repetition’ in the resulting proofs; in such cases, offer a streamlined proof using only basic rules. (This will give you a sense, both of the power of derived rules, and of how all the rules interact.)

Chapter 7

Natural deduction for Quantifier

31 Basic rules for Quantifier

Quantifier makes use of all of the connectives of Sentential. So proofs in Quantifier will use all of the basic and derived rules from chapter 6. We shall also use the proof-theoretic notions (particularly, the symbol ‘ \vdash ’) introduced in that chapter. However, we will also need some new basic rules to govern the quantifiers, and to govern the identity sign.

31.1 Universal elimination

Holding fixed the claim that everything is F, you can conclude that any particular thing is F. You name it; it’s F. So the following should be fine:

$$\begin{array}{l|l} 1 & \forall x Rxxd \\ \hline 2 & Raad \quad \forall E 1 \end{array}$$

We obtained line 2 by dropping the universal quantifier and replacing every instance of ‘ x ’ with ‘ a ’. Equally, the following should be allowed:

$$\begin{array}{l|l} 1 & \forall x Rxxd \\ \hline 2 & Rddd \quad \forall E 1 \end{array}$$

We obtained line 2 here by dropping the universal quantifier and replacing every instance of ‘ x ’ with ‘ d ’. We could have done the same with any other name we wanted.

This motivates the UNIVERSAL ELIMINATION rule ($\forall E$), using the notion for uniform substitution we introduced in §22.3:

$\begin{array}{l l} m & \forall x \mathcal{A} \\ & \vdots \\ & \mathcal{A} _{c \sim x} \quad \forall E m \end{array}$

The intent of the rule is that you can obtain any *substitution instance* of a universally quantified formula: replace every occurrence of the quantified variable x in \mathcal{A} with any chosen name. Remember here that the expression ‘ c ’ is a metalanguage variable over names: you are not required to replace the variable x by the Quantifier name ‘ c ’, but you can select any name you like!

I should emphasise that (as with every elimination rule) you can only apply the $\forall E$ rule when the universal quantifier is the main connective. Thus the following is outright banned:

1	$\forall x Bx \rightarrow Bk$	
2	$Bb \rightarrow Bk$	naughtily attempting to invoke $\forall E$ 1

This is illegitimate, since ‘ $\forall x$ ’ is not the main connective in line 1. (If you need a reminder as to why this sort of inference should be banned, reread §16.)

31.2 Existential introduction

Given the assumption that some specific named thing is an F, you can conclude that something is an F: ‘Sylvester reads, so someone reads’ seems like a conclusive argument. So we ought to allow:

1	$Raad$	
2	$\exists x Raax$	$\exists I$ 1

Here, we have replaced the name ‘ d ’ with a variable ‘ x ’, and then existentially quantified over it. Equally, we would have allowed:

1	$Raad$	
2	$\exists x Rxxd$	$\exists I$ 1

Here we have replaced both instances of the name ‘ a ’ with a variable, and then existentially generalised.

There are some pitfalls with this description of what we have done. The following argument is invalid: ‘Someone loves Alice; so someone is such that someone loves themselves’. So we ought not to be able to conclude ‘ $\exists x \exists x Rxx$ ’ from ‘ $\exists x Rxa$ ’. Accordingly, our rule cannot be *replace a name by a variable, and stick a corresponding quantifier out the front* – since that would permit the proof of the invalid argument.

We take our cue from the $\forall E$ rule. This rule says: take a formula \mathcal{A} , then we can remove the quantifier and substitute an arbitrary name for some variable. The $\exists I$ rule is in some sense a mirror image of this rule: it allows us to move from a sentence with an arbitrary name to a quantified formula. So here is how we formulate our rule of EXISTENTIAL INTRODUCTION:

m	$\mathcal{A} _{c \sim x}$	
	\vdots	
	$\exists x \mathcal{A}$	$\exists I m$

So really we should think that the proof on the previous page should be thought of as concluding $\exists x Rxxd$ from $'Rxxd'|_{a \sim x}$.

If we have this rule, we cannot provide a proof of the invalid argument. For $'\exists x Rxa'$ is not a substitution instance of $'\exists x \exists x Rxx'$ – both instances of $'x'$ in $'Rxx'$ are bound by the second existential quantifier, so neither is free to be substituted. So the premise is not of the right form for the rule of $\exists I$ to apply.

On the other hand, this proof is correct:

1	Raa	
2	$\exists x Rax$	$\exists I 1$

Why? Because the assumption $'Raa'$ is in fact not only a substitution instance of $\exists x Rxx$, but also a substitution instance of $\exists x Rax$, since $'Rax'|_{a \sim x}$ is just $'Raa'$ too. So we can vindicate the intuitively correct argument 'Narcissus loves himself, so there is someone who loves Narcissus'.

As we just saw, applying this rule requires some skill in being able to recognise substitution instances. Thus the following is allowed:

1	$Raad$	
2	$\exists x Rxad$	$\exists I 1$
3	$\exists y \exists x Rxyd$	$\exists I 2$

This is okay, because $'Raad'$ can arise from substitution of $'a'$ for $'x'$ in $'Rxad'$, and $\exists x Rxad$ can arise from substitution of $'a'$ for $'y'$ in $\exists x Rxyd$. But this is banned:

1	$Raad$	
2	$\exists xRxad$	$\exists I$ 1
3	$\exists x\exists xRxxd$	naughtily attempting to invoke $\exists I$ 2

This is because ' $\exists xRxad$ ' is not a substitution instance of ' $\exists x\exists xRxxd$ ', since (again) both occurrences of ' x ' in ' $Rxxd$ ' are already bound and so not available for free substitution.

31.3 Empty domains

The following proof combines our two new rules for quantifiers:

1	$\forall xFx$	
2	Fa	$\forall E$ 1
3	$\exists xFx$	$\exists I$ 2

Could this be a bad proof? If anything exists at all, then certainly we can infer that something is F, from the fact that everything is F. But what if *nothing* exists at all? Then it is surely vacuously true that everything is F; however, it does not follow that something is F, for there is nothing to *be* F. So if we claim that, as a matter of logic alone, ' $\exists xFx$ ' follows from ' $\forall xFx$ ', then we are claiming that, as a matter of *logic alone*, there is something rather than nothing. This might strike us as a bit odd.

Actually, we are already committed to this oddity. In §15, we stipulated that domains in Quantifier must have at least one member. We then defined a logical truth (of Quantifier) as a sentence which is true in every interpretation. Since ' $\exists x x = x$ ' will be true in every interpretation, this *also* had the effect of stipulating that it is a matter of logic that there is something rather than nothing.

Since it is far from clear that logic should tell us that there must be something rather than nothing, we might well be cheating a bit here.

If we refuse to cheat, though, then we pay a high cost. Here are three things that we want to hold on to:

- › $\forall xFx \vdash Fa$: after all, that was $\forall E$.
- › $Fa \vdash \exists xFx$: after all, that was $\exists I$.
- › the ability to copy-and-paste proofs together: after all, reasoning works by putting lots of little steps together into rather big chains.

If we get what we want on all three counts, then we have to countenance that $\forall xFx \vdash \exists xFx$. So, if we get what we want on all three counts, the proof system alone tells us that there is something rather than nothing. And if we refuse to accept that, then we have to surrender one of the three things that we want to hold on to!

Before we start thinking about which to surrender,¹ we might want to ask how *much* of a cheat this is. Granted, it may make it harder to engage in theological debates about why there is something rather than nothing. But the rest of the time, we will get along just fine. So maybe we should just regard our proof system (and Quantifier, more generally) as having a very slightly limited purview. If we ever want to allow for the possibility of *nothing*, then we shall have to cast around for a more complicated proof system. But for as long as we are content to ignore that possibility, our proof system is perfectly in order. (As, similarly, is the stipulation that every domain must contain at least one object.)

31.4 Universal introduction

Suppose you had shown of each particular thing that it is F (and that there are no other things to consider). Then you would be justified in claiming that everything is F. This would motivate the following proof rule. If you had established each and every single substitution instance of ' $\forall xFx$ ', then you can infer ' $\forall xFx$ '.

Unfortunately, that rule would be utterly unusable. To establish each and every single substitution instance would require proving ' Fa ', ' Fb ', ..., ' Fj_2 ', ..., ' Fr_{79002} ', ..., and so on. Indeed, since there are infinitely many names in Quantifier, this process would never come to an end. So we could never apply that rule. We need to be a bit more cunning in coming up with our rule for introducing universal quantification.

Our cunning thought will be inspired by considering:

$$\forall xFx \therefore \forall yFy$$

This argument should *obviously* be valid. After all, alphabetical variation in choice of variables ought to be a matter of taste, and of no logical consequence. But how might our proof system reflect this? Suppose we begin a proof thus:

1	$\forall xFx$	
2	Fa	$\forall E$ 1

¹ Many will opt for restricting $\exists I$. If we permit an empty domain, we will also need 'empty names' – names without a referent. When the name c is empty, it seems problematic to conclude from ' c is F' that there is something which is F. (Does 'Santa Claus drives a flying sleigh' entail 'Someone drives a flying sleigh'?) But empty names are not cost-free; understanding how a name that doesn't name anything can have any meaning at all has vexed many philosophers and linguists.

We have proved ‘ Fa ’. And, of course, nothing stops us from using the same justification to prove ‘ Fb ’, ‘ Fc ’, ..., ‘ Fj_2 ’, ..., ‘ Fr_{79002} ’, ..., and so on until we run out of space, time, or patience. But reflecting on this, we see that this is a way to prove Fc , for any name c . And if we can do it for *any* thing, we should surely be able to say that ‘ F ’ is true of *everything*. This therefore justifies us in inferring ‘ $\forall yFy$ ’, thus:

1	$\forall xFx$	
2	Fa	$\forall E$ 1
3	$\forall yFy$	$\forall I$ 2

The crucial thought here is that ‘ a ’ was just some *arbitrary* name. There was nothing special about it – we might have chosen any other name – and still the proof would be fine. And this crucial thought motivates the universal introduction rule ($\forall I$):

m	$\mathcal{A} _{c \curvearrowright x}$	
	⋮	
	$\forall x\mathcal{A}$	$\forall I$ m
c must not occur in any undischarged assumption, or elsewhere in \mathcal{A}		

A crucial aspect of this rule, though, is bound up in the accompanying constraint. In English, a name like ‘Sylvester’ can play two roles: it can be introduced as a name for a specific thing (‘let me dub thee Sylvester!’), or as an *arbitrary name*, introduced by this sort of stipulation: *let ‘Sylvester’ name some arbitrarily chosen man*. The name doesn’t tell us, as it subsequently appears, whether it was introduced in one way or the other. But if it was introduced as an arbitrary name, then any conclusions we draw about this Sylvester aren’t really dependent on the particular arbitrarily chosen referent – they all depend rather on the stipulation used in introducing the name, and so (specifically) they will all be consequences of the only fact we know for sure about this Sylvester, that he is male. If all men are mortal, then an arbitrarily chosen man, whom we temporarily call ‘Sylvester’, is mortal. If Sylvester is mortal, then there is a date he will die. But since he was selected arbitrarily, without reference to any further particulars of his life, then for any man, there exists a date he will die. And that is appropriate reasoning from a universal generalisation, to another generalisation, via claims about a specific but arbitrarily chosen person.²

We don’t have stipulations to introduce a name as an arbitrary name in Quantifier. But we do have a way of ensuring that the name has no prior associations other than those

² The details about how this sort of arbitrary reference works are interesting. A controversial but nevertheless attractive view of how it might work is Wylie Breckenridge and Ofra Magidor, ‘Arbitrary Reference’, *Philosophical Studies* 158 (2012), pp. 377–400.

linked to a prior universal generalisation, if we insist that, when the name is about to be eliminated from the proof, no assumption about what that name denotes is being relied on. That way, we can know that however it was introduced to the proof, it was not done in a way that involved making specific assumptions about whatever the name arbitrarily picks out. The simplest way for this to happen is that the name was introduced by an application of $\forall E$, as an arbitrary name in the standard sense. But there are other ways too.

This constraint ensures that we are always reasoning at a sufficiently general level. To see the constraint in action, consider this terrible argument:

Everyone loves Kylie Minogue; therefore everyone loves themselves.

We might symbolise this obviously invalid inference pattern as:

$$\forall xLxk \therefore \forall xLxx$$

Now, suppose we tried to offer a proof that vindicates this argument:

1	$\forall xLxk$	
2	Lkk	$\forall E$ 1
3	$\forall xLxx$	naughtily attempting to invoke $\forall I$ 2

This is not allowed, because ‘ k ’ occurred already in an undischarged assumption, namely, on line 1. The crucial point is that, if we have made any assumptions about the object we are working with (including assumptions embedded in \mathcal{A} itself), then we are not reasoning generally enough to license the use of $\forall I$.

Although the name may not occur in any *undischarged* assumption, it may occur as a discharged assumption. That is, it may occur in a subproof that we have already closed. For example:

1	Gd	
2	Gd	R 1
3	$Gd \rightarrow Gd$	$\rightarrow I$ 1-2
4	$\forall z(Gz \rightarrow Gz)$	$\forall I$ 3

This tells us that ‘ $\forall z(Gz \rightarrow Gz)$ ’ is a *theorem*. And that is as it should be.

Here is another proof featuring an application of $\forall I$ after discharging an assumption about some name ‘ a ’:

1	$Fa \wedge \neg Fa$	
2	Fa	$\rightarrow E$ 1
3	$\neg Fa$	$\rightarrow E$ 1
4	$\neg(Fa \wedge \neg Fa)$	$\neg I$ 1-2, 1-3
5	$\forall x \neg(Fx \wedge \neg Fx)$	$\forall I$ 4

Here we were able to derive that something could not be true of a , no matter what a is. So that property could not be true of anything at all. We cannot even coherently assume this about a , so it doesn't really matter what ' a ' denotes, it won't fulfil the assumption made about it. That is why we are entitled to discharge that assumption, and then any subsequent use of ' a ' in the proof must be depending not on particular facts about this a , but about anything at all, including whatever it is that ' a ' happens to pick out.

31.5 Existential elimination

Suppose we know that *something* is F. The problem is that simply knowing this does not tell us which thing is F. So it would seem that from ' $\exists xFx$ ' we cannot immediately conclude ' Fa ', ' Fe_{23} ', or any other substitution instance of the sentence. What can we do?

Suppose we know that something is F, and that everything which is F is G. In (almost) natural English, we might reason thus:

Since something is F, there is some particular thing which is an F. We do not know anything about it, other than that it's an F, but for convenience, let's call it 'obbie'. So: obbie is F. Since everything which is F is G, it follows that obbie is G. But since obbie is G, it follows that something is G. And nothing depended on which object, exactly, obbie was. So, something is G.

We might try to capture this reasoning pattern in a proof as follows:

1	$\exists xFx$	
2	$\forall x(Fx \rightarrow Gx)$	
3	Fo	
4	$Fo \rightarrow Go$	$\forall E$ 2
5	Go	$\rightarrow E$ 4, 3
6	$\exists xGx$	$\exists I$ 5
7	$\exists xGx$	$\exists E$ 1, 3-6

Breaking this down: we started by writing down our assumptions. At line 3, we made an additional assumption: ' Fo '. This was just a substitution instance of ' $\exists xFx$ '. On this assumption, we established ' $\exists xGx$ '. But note that we had made no *special* assumptions about the object named by ' o '; we had *only* assumed that it satisfies ' Fx '. So nothing depends upon which object it is. And line 1 told us that *something* satisfies ' Fx '. So our reasoning pattern was perfectly general. We can discharge the specific assumption ' Fo ', and simply infer ' $\exists xGx$ ' on its own.

Putting this together, we obtain the existential elimination rule ($\exists E$):

m	$\exists xA$				
i	<table style="border-left: 1px solid black; padding-left: 10px;"> <tr> <td style="border-bottom: 1px solid black;">$A _{c \sim x}$</td> </tr> <tr> <td>\vdots</td> </tr> <tr> <td>B</td> </tr> </table>	$A _{c \sim x}$	\vdots	B	
$A _{c \sim x}$					
\vdots					
B					
j	B				
	\vdots				
	B	$\exists E\ m, i-j$			

c must not occur in any assumption undischarged before line i
 c must not occur in $\exists xA$
 c must not occur in B

As with universal introduction, the constraints are extremely important. To see why, consider the following terrible argument:

Tim Button is a lecturer. There is someone who is not a lecturer. So Tim Button is both a lecturer and not a lecturer.

We might symbolise this obviously invalid inference pattern as follows:

$$Lb, \exists x \neg Lx \therefore Lb \wedge \neg Lb$$

Now, suppose we tried to offer a proof that vindicates this argument:

1	Lb			
2	$\exists x \neg Lx$			
	<table style="border-left: 1px solid black; padding-left: 10px;"> <tr> <td style="border-bottom: 1px solid black;">$\neg Lb$</td> </tr> <tr> <td>$Lb \wedge \neg Lb$</td> </tr> </table>	$\neg Lb$	$Lb \wedge \neg Lb$	$\wedge E\ 1, 3$
$\neg Lb$				
$Lb \wedge \neg Lb$				
4	$Lb \wedge \neg Lb$			
5	$Lb \wedge \neg Lb$	naughtily attempting to invoke $\exists E\ 2, 3-4$		

The last line of the proof is not allowed. The name that we used in our substitution instance for ' $\exists x\neg Lx$ ' on line 3, namely ' b ', occurs in line 4. And the following proof would be no better:

1	Lb	
2	$\exists x\neg Lx$	
3	$\neg Lb$	
4	$Lb \wedge \neg Lb$	$\wedge E$ 1, 3
5	$\exists x(Lx \wedge \neg Lx)$	$\exists I$ 4
6	$\exists x(Lx \wedge \neg Lx)$	naughtily attempting to invoke $\exists E$ 2, 3–5

The last line of the proof would still not be allowed. For the name that we used in our substitution instance for ' $\exists x\neg Lx$ ', namely ' b ', occurs in an undischarged assumption, namely line 1.

The moral of the story is this.

If you want to squeeze information out of an existentially quantified claim $\exists x\mathcal{A}$, choose a *new* name, never before used in the proof, to substitute for the variable in \mathcal{A} .

That way, you can guarantee that you meet all the constraints on the rule for $\exists E$.

An argument that makes use of both patterns of arbitrary reasoning is this example due to Breckenridge and Magidor: 'from the premise that there is someone who loves everyone to the conclusion that everyone is such that someone loves them'. Here is a proof in our system, letting ' Lxy ' symbolise ' $\text{---}x$ loves $\text{---}y$ ', letting ' h ' symbolise the arbitrarily chosen name 'Hiccup' and letting ' a ' symbolise the arbitrarily chosen name 'Astrid':

1	$\exists x\forall yLxy$	
2	$\forall yLhy$	
3	Lha	$\forall E$ 2
4	$\exists xLxa$	$\exists I$ 3
5	$\forall y\exists xLxy$	$\forall I$ 4
6	$\forall y\exists xLxy$	$\exists E$ 1, 2–5

At line 3, both our arbitrary names are in play – h was newly introduced to the proof in line 2 as the arbitrary person Hiccup who witnesses the truth of ‘ $\exists x\forall yLxy$ ’, and ‘ a ’ at line 3 as an arbitrary person Astrid beloved by Hiccup. We can apply $\exists I$ without restriction at line 4, which takes the name ‘ h ’ out of the picture – we no longer rely on the specific instance chosen, since we are back at generalities about someone who loves everyone, being such that they also love the arbitrarily chosen someone Astrid. So we can safely apply $\forall I$ at line 5, since the name ‘ a ’ appears in no assumption nor in ‘ $\forall y\exists xLxy$ ’. But now we have at line 5 a claim that doesn’t involve the arbitrary name ‘ h ’ either, which was newly chosen to not be in any undischarged assumption or in $\exists x\forall yLxy$. So we can safely say that the name Hiccup was just arbitrary, and nothing in the proof of ‘ $\forall y\exists xLxy$ ’ depended on it, so we can discharge the specific assumption about h that was used in the course of that proof and nevertheless retain our entitlement to ‘ $\forall y\exists xLxy$ ’.

31.6 Justification of these quantifier rules

Above, I offered informal arguments for each of our quantifier rules that seem to exemplify the pattern of argument in the rule, and to be intuitively valid. But we can also offer justifications for our rules in terms of interpretations of the sentences involved, and the principles governing truth of quantified sentences introduced in §22.3.

For example, consider any interpretation which makes $\forall x\mathcal{A}$ true. In any such interpretation, there will be a non-empty domain, and every name will denote some member of this domain. $\forall x\mathcal{A}$ is true just in case for any name we like, it will denote something of which $\mathcal{A}|_{c\rightsquigarrow x}$ is true. So in any such interpretation, for each name in the language c , $\mathcal{A}|_{c\rightsquigarrow x}$ will also be true. So the proof rule of $\forall E$ corresponds to a valid argument form.

For $\exists E$, the case is only a little more involved. Suppose $\exists x\mathcal{A}$ is true in an interpretation. Then there is some interpretation, otherwise just like the original one, in which some new name c is assigned to some object in the domain, and where $\mathcal{A}|_{c\rightsquigarrow x}$ is true. Suppose that, in fact, every interpretation which makes $\mathcal{A}|_{c\rightsquigarrow x}$ true also makes \mathcal{B} true, where the new name c does not appear in \mathcal{B} . Could \mathcal{B} be false in our original interpretation? No – for everything that appears in \mathcal{B} is already interpreted in the original interpretation, with the same interpretation as in the interpretation which makes it true. So it must be true in our original interpretation too. So $\mathcal{C}_1, \dots, \mathcal{C}_n, \exists x\mathcal{A} \models \mathcal{B}$ (when the name c makes no appearance in any sentence in this argument), and the proof rule of $\exists E$ corresponds to a valid argument form.

You may also offer arguments from interpretations to the effect that our other quantifier proof rules correspond to valid arguments in Quantifier:

- › $\mathcal{A}|_{c\rightsquigarrow x} \models \exists x\mathcal{A}$;
- › $\mathcal{C}_1, \dots, \mathcal{C}_n, \mathcal{A}|_{c\rightsquigarrow x} \models \exists x\mathcal{A}$, so long as c appears nowhere in the argument except in $\mathcal{A}|_{c\rightsquigarrow x}$.

So we are again comforted: our proof rules can never lead us from true assumptions to false claims, if correctly applied.

Practice exercises

A. The following two ‘proofs’ are *incorrect*. Explain why both are incorrect. Also, provide interpretations which would invalidate the fallacious argument forms the ‘proofs’ enshrine:

1	$\forall xRxx$	
2	Raa	$\forall E$ 1
3	$\forall yRay$	$\forall I$ 2
4	$\forall x\forall yRxy$	$\forall I$ 3

1	$\forall x\exists yRxy$					
2	$\exists yRay$	$\forall E$ 1				
3	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">Raa</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\exists xRxx$</td> <td style="padding-left: 10px;">$\exists I$ 3</td> </tr> </table>	Raa		$\exists xRxx$	$\exists I$ 3	
Raa						
$\exists xRxx$	$\exists I$ 3					
5	$\exists xRxx$	$\exists E$ 2, 3-4				

B. The following three proofs are missing their citations (rule and line numbers). Add them, to turn them into bona fide proofs.

1	$\forall x\exists y(Rxy \vee Ryx)$			
2	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\forall x\neg Rmx$</td> <td></td> </tr> </table>	$\forall x\neg Rmx$		
$\forall x\neg Rmx$				
3	$\exists y(Rmy \vee Rym)$			
4	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$Rma \vee Ram$</td> <td></td> </tr> </table>	$Rma \vee Ram$		
$Rma \vee Ram$				
5	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\neg Rma$</td> <td></td> </tr> </table>	$\neg Rma$		
$\neg Rma$				
6	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">Ram</td> <td></td> </tr> </table>	Ram		
Ram				
7	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">$\exists xRxm$</td> <td></td> </tr> </table>	$\exists xRxm$		
$\exists xRxm$				
8	$\exists xRxm$			

1	$\forall x(\exists yLxy \rightarrow \forall zLzx)$			
2	<table style="border-collapse: collapse; margin-left: 10px;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;">Lab</td> <td></td> </tr> </table>	Lab		
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	$\forall x(Jx \rightarrow Kx)$
2	$\exists x\forall yLxy$
3	$\forall xJx$
4	$\forall yLay$
5	Laa
6	Ja
7	$Ja \rightarrow Ka$
8	Ka
9	$Ka \wedge Laa$
10	$\exists x(Kx \wedge Lxx)$
11	$\exists x(Kx \wedge Lxx)$

C. In §16 problem part A, we considered fifteen syllogistic figures of Aristotelian logic. Provide proofs for each of the argument forms. NB: You will find it *much* easier if you symbolise (for example) ‘No F is G’ as ‘ $\forall x(Fx \rightarrow \neg Gx)$ ’.

D. Aristotle and his successors identified other syllogistic forms which depended upon ‘existential import’. Symbolise each of the following argument forms in *Quantifier* and offer proofs.

- › **Barbari**. Something is H. All G are F. All H are G. So: Some H is F
- › **Celaront**. Something is H. No G are F. All H are G. So: Some H is not F
- › **Cesaro**. Something is H. No F are G. All H are G. So: Some H is not F.
- › **Camestros**. Something is H. All F are G. No H are G. So: Some H is not F.
- › **Felapton**. Something is G. No G are F. All G are H. So: Some H is not F.
- › **Darapti**. Something is G. All G are F. All G are H. So: Some H is F.
- › **Calemos**. Something is H. All F are G. No G are H. So: Some H is not F.
- › **Fesapo**. Something is G. No F is G. All G are H. So: Some H is not F.
- › **Bamalip**. Something is F. All F are G. All G are H. So: Some H are F.

E. Provide a proof of each claim.

1. $\vdash \forall xFx \vee \neg\forall xFx$
2. $\vdash \forall z(Pz \vee \neg Pz)$
3. $\forall x(Ax \rightarrow Bx), \exists xAx \vdash \exists xBx$
4. $\forall x(Mx \leftrightarrow Nx), Ma \wedge \exists xRxa \vdash \exists xNx$
5. $\forall x\forall yGxy \vdash \exists xGxx$
6. $\vdash \forall xRxx \rightarrow \exists x\exists yRxy$
7. $\vdash \forall y\exists x(Qy \rightarrow Qx)$
8. $Na \rightarrow \forall x(Mx \leftrightarrow Ma), Ma, \neg Mb \vdash \neg Na$
9. $\forall x\forall y(Gxy \rightarrow Gyx) \vdash \forall x\forall y(Gxy \leftrightarrow Gyx)$
10. $\forall x(\neg Mx \vee Ljx), \forall x(Bx \rightarrow Ljx), \forall x(Mx \vee Bx) \vdash \forall xLjx$

F. Write a symbolisation key for the following argument, symbolise it, and prove it:

There is someone who likes everyone who likes everyone that she likes. Therefore, there is someone who likes herself.

G. For each of the following pairs of sentences: If they are provably equivalent, give proofs to show this. If they are not, construct an interpretation to show that they are not logically equivalent.

1. $\forall xPx \rightarrow Qc, \forall x(Px \rightarrow Qc)$
2. $\forall x\forall y\forall zBxyz, \forall xBxxx$
3. $\forall x\forall yDxy, \forall y\forall xDxy$
4. $\exists x\forall yDxy, \forall y\exists xDxy$
5. $\forall x(Rca \leftrightarrow Rxa), Rca \leftrightarrow \forall xRxa$

H. For each of the following arguments: If it is valid in Quantifier, give a proof. If it is invalid, construct an interpretation to show that it is invalid.

1. $\exists y\forall xRxy \therefore \forall x\exists yRxy$
2. $\exists x(Px \wedge \neg Qx) \therefore \forall x(Px \rightarrow \neg Qx)$
3. $\forall x(Sx \rightarrow Ta), Sd \therefore Ta$
4. $\forall x(Ax \rightarrow Bx), \forall x(Bx \rightarrow Cx) \therefore \forall x(Ax \rightarrow Cx)$
5. $\exists x(Dx \vee Ex), \forall x(Dx \rightarrow Fx) \therefore \exists x(Dx \wedge Fx)$
6. $\forall x\forall y(Rxy \vee Ryx) \therefore Rjj$
7. $\exists x\exists y(Rxy \vee Ryx) \therefore Rjj$
8. $\forall xPx \rightarrow \forall xQx, \exists x\neg Px \therefore \exists x\neg Qx$

32 Derived rules in Quantifier: conversion of quantifiers

In this section, we shall add some additional rules to the basic rules of the previous section. These govern the interaction of quantifiers and negation. But they are no substantive addition to our basic rules: for each of the proposed additions, it can be shown that their role in any proof can be wholly emulated by some suitable applications of our basic rules from §31. (The point here is as in §30.)

In §15, we noted that $\neg\exists x\mathcal{A}$ is logically equivalent to $\forall x\neg\mathcal{A}$. We shall add some rules to our proof system that govern this. In particular, we add two rules, one for each direction of the equivalence:

$\frac{m \quad \forall x\neg\mathcal{A}}{\neg\exists x\mathcal{A}} \quad \text{CQ } m$	$\frac{m \quad \neg\exists x\mathcal{A}}{\forall x\neg\mathcal{A}} \quad \text{CQ } m$
--	--

Here is a schematic proof corresponding to our first CQ rule:

1	$\forall x\neg\mathcal{A}$															
2	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\exists x\mathcal{A}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\mathcal{A} _{c \rightsquigarrow x}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table> </td> <td style="padding-left: 10px;">$\neg I$ 4-5, 4-3</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$B \wedge \neg B$</td> <td style="padding-left: 10px;">$\neg E$ 6</td> </tr> </table> </td> <td style="padding-left: 10px;">$\exists E$ 2, 3-7</td> </tr> </table>	$\exists x\mathcal{A}$		<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\mathcal{A} _{c \rightsquigarrow x}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table> </td> <td style="padding-left: 10px;">$\neg I$ 4-5, 4-3</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$B \wedge \neg B$</td> <td style="padding-left: 10px;">$\neg E$ 6</td> </tr> </table>	$\mathcal{A} _{c \rightsquigarrow x}$		<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table>	$\neg(B \wedge \neg B)$		$\neg\mathcal{A} _{c \rightsquigarrow x}$	$\forall E$ 1	$\neg I$ 4-5, 4-3	$B \wedge \neg B$	$\neg E$ 6	$\exists E$ 2, 3-7	
$\exists x\mathcal{A}$																
<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\mathcal{A} _{c \rightsquigarrow x}$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table> </td> <td style="padding-left: 10px;">$\neg I$ 4-5, 4-3</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$B \wedge \neg B$</td> <td style="padding-left: 10px;">$\neg E$ 6</td> </tr> </table>	$\mathcal{A} _{c \rightsquigarrow x}$		<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table>	$\neg(B \wedge \neg B)$		$\neg\mathcal{A} _{c \rightsquigarrow x}$	$\forall E$ 1	$\neg I$ 4-5, 4-3	$B \wedge \neg B$	$\neg E$ 6	$\exists E$ 2, 3-7					
$\mathcal{A} _{c \rightsquigarrow x}$																
<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg(B \wedge \neg B)$</td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">$\neg\mathcal{A} _{c \rightsquigarrow x}$</td> <td style="padding-left: 10px;">$\forall E$ 1</td> </tr> </table>	$\neg(B \wedge \neg B)$		$\neg\mathcal{A} _{c \rightsquigarrow x}$	$\forall E$ 1	$\neg I$ 4-5, 4-3											
$\neg(B \wedge \neg B)$																
$\neg\mathcal{A} _{c \rightsquigarrow x}$	$\forall E$ 1															
$B \wedge \neg B$	$\neg E$ 6															
9	$\neg\exists x\mathcal{A}$	$\neg I$ 2-8														

A couple of things to note about this proof.

1. I was hasty at line 9 – officially I ought to have applied $\wedge E$ to line 8, obtaining the contradictory conjuncts in the subproof, and then applied $\neg I$ to the assumption opening that subproof. (But then the proof would have gone over the page.)
2. Note that we had to introduce the new name c at line 3. Once we did so, there was no obstacle to applying $\forall E$ on that newly introduced name in line 5. But if we had done things the other way around, applying $\forall E$ first to some new name c , we would have had to open the subproof with yet another new name d .
3. The sentence B cannot contain the name c if the application of $\exists E$ at line 8 is to be correct. We introduce this arbitrary contradiction precisely so we can show that the contradictoriness of our initial assumptions does not depend on the particular choice of name. The alternative would have been to show that the assumption $\mathcal{A}|_{c \sim x}$ leads to contradiction, and then applied $\neg I$ – but that would have left the name c outside the scope of a subproof and would not have allowed us to apply $\exists E$.

Here is a schematic basic proof emulating the second CQ rule, in which some of the same issues arise:

1	$\exists x \neg Ax$	
2	$\forall x Ax$	
3	$\neg Ac$	
4	$\neg(B \wedge \neg B)$	
5	Ac	$\forall E$ 2
6	$\neg\neg(B \wedge \neg B)$	$\neg I$ 4-3, 4-5
7	$B \wedge \neg B$	$\neg E$ 6
8	$B \wedge \neg B$	$\exists E$ 1, 3-7
9	$\neg\forall x Ax$	$\neg I$ 2-8

Equally, we add rules corresponding to the equivalence of $\exists x \neg \mathcal{A}$ and $\neg\forall x \mathcal{A}$:

m	$\exists x \neg \mathcal{A}$		m	$\neg\forall x \mathcal{A}$	
	$\neg\forall x \mathcal{A}$	$CQ\ m$		$\exists x \neg \mathcal{A}$	$CQ\ m$

Similar schematic proofs can be offered for the other two CQ rules that show they too can be emulated using just our basic rules.

Practice exercises

A. Show that the following are jointly contrary:

1. $Sa \rightarrow Tm, Tm \rightarrow Sa, Tm \wedge \neg Sa$
2. $\neg \exists x Rxa, \forall x \forall y Ryx$
3. $\neg \exists x \exists y Lxy, Laa$
4. $\forall x (Px \rightarrow Qx), \forall z (Pz \rightarrow Rz), \forall y Py, \neg Qa \wedge \neg Rb$

B. Show that each pair of sentences is provably equivalent:

1. $\forall x (Ax \rightarrow \neg Bx), \neg \exists x (Ax \wedge Bx)$
2. $\forall x (\neg Ax \rightarrow Bx), \forall x Ax \vee Bx$

C. In §16, I considered what happens when we move quantifiers ‘across’ various connectives. Show that each pair of sentences is provably equivalent:

1. $\forall x (Fx \wedge Ga), \forall x Fx \wedge Ga$
2. $\exists x (Fx \vee Ga), \exists x Fx \vee Ga$
3. $\forall x (Ga \rightarrow Fx), Ga \rightarrow \forall x Fx$
4. $\forall x (Fx \rightarrow Ga), \exists x Fx \rightarrow Ga$
5. $\exists x (Ga \rightarrow Fx), Ga \rightarrow \exists x Fx$
6. $\exists x (Fx \rightarrow Ga), \forall x Fx \rightarrow Ga$

NB: the variable ‘ x ’ does not occur in ‘ Ga ’.

When all the quantifiers occur at the beginning of a sentence, that sentence is said to be in *prenex normal form*. These equivalences are sometimes called *prenexing rules*, since they give us a means for putting any sentence into prenex normal form.

D. Offer proofs which justify the addition of the third and fourth CQ rules as derived rules.

33 Rules for identity

In §21, I mentioned the philosophically contentious thesis of the *identity of indiscernibles*. This is the claim that objects which are indiscernible in every way – which means, for us, that exactly the same predicates are true of both objects – are, in fact, identical to each other. I also mentioned that in *Quantifier*, this thesis is not true. There are interpretations in which, for each property denoted by any predicate of the language, two distinct objects either both have that property, or both lack that property. They may well differ on some property ‘in reality’, but there is nothing in *Quantifier* which guarantees that that property is assigned as the interpretation of any predicate of the language. It follows that, no matter how many sentences of *Quantifier* about those two objects I assume, those sentences will not entail that these distinct objects are identical (luckily for us). Unless, of course, you tell me that the two objects are, in fact, identical. But then the argument will hardly be very illuminating.

The consequence of this, for our proof system, is that 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 names.

However, every object is identical to itself. No premises, then, are required in order to conclude that something is identical to itself. So this will be the identity introduction rule:

$$\boxed{\quad | \quad c = c \quad =I}$$

Notice that this rule does not require referring to any prior lines of the proof, nor does it rely on any assumptions. For any name c , you can write $c = c$ at any point, with only the =I rule as justification.

Our elimination rule is more fun. If you have established ‘ $a = b$ ’, then anything that is true of the object named by ‘ a ’ must also be true of the object named by ‘ b ’, since it just is the object named by ‘ a ’. So for any sentence with ‘ a ’ in it, given the prior claim that ‘ $a = b$ ’, you can replace some or all of the occurrences of ‘ a ’ with ‘ b ’ and produce an equivalent sentence. For example, from ‘ Raa ’ and ‘ $a = b$ ’, you are justified in inferring ‘ Rab ’, ‘ Rba ’ or ‘ Rbb ’. More generally:

m	$a = b$	
n	$\mathcal{A} _{a \sim x}$	
	$\mathcal{A} _{b \sim x}$	$=E\ m, n$

This uses our standard notion of substitution – it basically says that if you have some sentence which arises from substituting ‘ a ’ for some variable in a formula, then you are entitled to another substitution instance of the same formula using ‘ b ’ instead. Lines m and n can occur in either order, and do not need to be adjacent, but we always cite the statement of identity first.

Note that nothing in the rule forbids the constant ‘ b ’ from occurring in \mathcal{A} . So this is a perfectly good instance of the rule:

1	$a = b$	
2	Rab	
3	Rbb	$=E\ 1, 2$

Here, ‘ Rab ’ is ‘ Rxb ’ $|_{a \sim x}$, and the conclusion ‘ Rbb ’ is ‘ Rxb ’ $|_{b \sim x}$, which conforms to the rule. This formulation allows us, in effect, to replace some-but-not-all occurrences of a name in a sentence by a co-referring name. This rule is sometimes called *Leibniz’s Law*, after Gottfried Leibniz.

To see the rules in action, we shall prove some quick results. First, we shall prove that identity is *symmetric*:

1	$a = b$	
2	$a = a$	$=I$
3	$b = a$	$=E\ 1, 2$
4	$a = b \rightarrow b = a$	$\rightarrow I\ 1-3$
5	$\forall y(a = y \rightarrow y = a)$	$\forall I\ 4$
6	$\forall x \forall y(x = y \rightarrow y = x)$	$\forall I\ 5$

Line 2 is just ‘ $x = a$ ’ $|_{a \sim x}$, as well as being of the right form for $=I$, and line 3 is just ‘ $x = a$ ’ $|_{b \sim x}$, so the move from 2 to 3 is in conformity with the $=E$ rule given the opening assumption.

Having noted the symmetry of identity, note that we can use this to establish the following schematic proof that allows us to use $a = b$ to also move from a claim about b to a claim about a , not just *vice versa* as in our =E rule.:

m	$a = b$	
$m + 1$	$a = a$	=I
$m + 2$	$b = a$	=E $m, m + 1$
n	$\mathcal{A} _{b \sim x}$	
	$\mathcal{A} _{a \sim x}$	=E $m + 2, n$

This schematic proof actually justifies this derived rule, to save time:

m	$a = b$	
n	$\mathcal{A} _{b \sim x}$	
	$\mathcal{A} _{a \sim x}$	=ES m, n

Second, we shall prove that identity is *transitive*:

1	$a = b \wedge b = c$	
2	$b = c$	∧E 1
3	$a = b$	∧E 1
4	$a = c$	=E 2, 3
5	$(a = b \wedge b = c) \rightarrow a = c$	→I 1-4
6	$\forall z((a = b \wedge b = z) \rightarrow a = z)$	∀I 5
7	$\forall y \forall z((a = y \wedge y = z) \rightarrow a = z)$	∀I 6
8	$\forall x \forall y \forall z((x = y \wedge y = z) \rightarrow x = z)$	∀I 7

We obtain line 4 by replacing 'b' in line 3 with 'c'; this is justified given line 2, ' $b = c$ '. We could alternatively have used the derived rule =ES to replace 'b' in line 2 with 'a', justified by line 3, ' $a = b$ '.

Here is a final proof, that identity is *reflexive*:

1	$a = a$	eI
2	$\forall x x = x$	$\forall I$ 1

This seems like magic! But note that the first line is not an assumption (there is no horizontal line), and hence not an undischarged assumption. So the constant ‘ a ’ appears in no undischarged assumption or anywhere in the proof other than in ‘ $x = x$ ’ _{$a \sim x$} , so the conclusion ‘ $\forall x x = x$ ’ follows by legitimate application of the $\forall I$ rule.

Practice exercises

A. 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. $\forall x x = m, Rma \vdash \exists x Rxx$
4. $\forall x \forall y (Rxy \rightarrow x = y) \vdash Rab \rightarrow Rba$
5. $\neg \exists x \neg x = m \vdash \forall x \forall y (Px \rightarrow Py)$
6. $\exists x Jx, \exists x \neg Jx \vdash \exists x \exists y \neg x = y$
7. $\forall x (x = n \leftrightarrow Mx), \forall x (Ox \vee \neg Mx) \vdash On$
8. $\exists x Dx, \forall x (x = p \leftrightarrow Dx) \vdash Dp$
9. $\exists x ((Kx \wedge \forall y (Ky \rightarrow x = y)) \wedge Bx), Kd \vdash Bd$
10. $\vdash Pa \rightarrow \forall x (Px \vee \neg x = a)$

B. Show that the following are provably equivalent:

- › $\exists x ((Fx \wedge \forall y (Fy \rightarrow x = y)) \wedge x = n)$
- › $Fn \wedge \forall y (Fy \rightarrow n = y)$

And hence that both have a decent claim to symbolise the English sentence ‘Nick is the F’.

C. In §18, I claimed that the following are logically equivalent symbolisations of the English sentence ‘there is exactly one F’:

- › $\exists x Fx \wedge \forall x \forall y ((Fx \wedge Fy) \rightarrow x = y)$
- › $\exists x (Fx \wedge \forall y (Fy \rightarrow x = y))$
- › $\exists x \forall y (Fy \leftrightarrow x = y)$

Show that they are all provably equivalent. (*Hint*: to show that three claims are provably equivalent, it suffices to show that the first proves the second, the second proves the third and the third proves the first; think about why.)

D. Symbolise the following argument

There is exactly one F. There is exactly one G. Nothing is both F and G. So:
there are exactly two things that are either F or G.

And offer a proof of it.

34 Proof-theoretic concepts and semantic concepts

We have used two different turnstiles in this book. This:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash \mathcal{C}$$

means that there is some proof which starts with assumptions $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ and ends with \mathcal{C} (and no undischarged assumptions other than $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$). This is a *proof-theoretic notion*.

By contrast, this:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vDash \mathcal{C}$$

means that there is no valuation (or interpretation) which makes all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ true and makes \mathcal{C} false. This concerns assignments of truth and falsity to sentences. It is a *semantic notion*.

I cannot emphasise enough that these are different notions. But I can emphasise it a bit more: *They are different notions*.

Once you have internalised this point, continue reading.

Although our semantic and proof-theoretic notions are different, there is a deep connection between them. To explain this connection, I shall start by considering the relationship between logical truths and theorems.

To show that a sentence is a theorem, you need only perform a proof. Granted, 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 logical truth, though, requires reasoning about all possible interpretations. Given a choice between showing that a sentence is a theorem and showing that it is a logical truth, it would be easier to show that it is a theorem.

Contrawise, to show that a sentence is *not* a theorem is hard. We would need to reason about all (possible) proofs. That is very difficult. But to show that a sentence is not a logical truth, you need only construct an interpretation in which the sentence is false. Granted, it

may be hard to come up with the interpretation; but once you have done so, it is relatively straightforward to check what truth value it assigns to a sentence. Given a choice between showing that a sentence is not a theorem and showing that it is not a logical truth, it would be easier to show that it is not a logical truth.

Fortunately, *a sentence is a theorem if and only if it is a logical truth*. As a result, if we provide a proof of \mathcal{A} on no assumptions, and thus show that \mathcal{A} is a theorem, we can legitimately infer that \mathcal{A} is a logical truth; i.e., $\models \mathcal{A}$. Similarly, if we construct an interpretation in which \mathcal{A} is false and thus show that it is not a logical truth, it follows that \mathcal{A} is not a theorem.

More generally, we have the following powerful result:

$$\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash \mathcal{B} \text{ iff } \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \models \mathcal{B}.$$

This shows that, whilst provability and entailment are *different* notions, they are extensionally equivalent, holding between just the same sentences in our languages. As such:

- › An argument is *valid* iff *the conclusion can be proved from the premises*.
- › Two sentences are *logically equivalent* iff *they are provably equivalent*.
- › Sentences are *jointly consistent* iff *they are not jointly contrary*.

For this reason, you can pick and choose when to think in terms of proofs and when to think in terms of valuations/interpretations, doing whichever is easier for a given task. The table on the next page summarises which is (usually) easier.

It is intuitive that provability and semantic entailment should agree. But – let me repeat this – do not be fooled by the similarity of the symbols ‘ \models ’ and ‘ \vdash ’. These two symbols have very different meanings. And the fact that provability and semantic entailment agree is not an easy result to come by.

We showed part of this result along the way, actually. All those little observations I made about how our proof rules were good each took the form of an argument that whenever $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash \mathcal{B}$, then also $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \models \mathcal{B}$. This is sometimes known as the **SOUNDNESS** of our proof system, that whenever you can prove something, there is a corresponding valid argument.

The other half of this result is more demanding. This is the **COMPLETENESS** of our proof system, that whenever there is a valid argument, there is a proof. It is somewhere around this point, of demonstrating the completeness half of the agreement between provability and semantic entailment, when introductory logic becomes intermediate logic. More than one of the authors of this book have written distinct sequels in which the completeness of *Sentential* is established, along with other results of a ‘metalogical’ flavour. The interested student is directed to these texts:

- › Tim Button's book is *Metatheory*; it covers just Sentential, and some of the results there were actually covered in this book already, notably results around expressiveness of Sentential in §14. *Metatheory* is available at people.ds.cam.ac.uk/tecb2/metatheory.shtml. It should be accessible for self-study by students who have successfully completed this course.
- › Antony Eagle's book *Elements of Deductive Logic* goes rather further than *Metatheory*, including direct proofs of Compactness for Sentential, consideration of alternative derivation systems, discussion of computation and decidability, and metatheory for Quantifier. It is available at github.com/antonyeagle/edl.

	Yes	No
Is \mathcal{A} a logical truth ?	give a proof which shows $\vdash \mathcal{A}$	give an interpretation in which \mathcal{A} is false
Is \mathcal{A} a contradiction ?	give a proof which shows $\vdash \neg \mathcal{A}$	give an interpretation in which \mathcal{A} is true
Are \mathcal{A} and \mathcal{B} equivalent ?	give two proofs, one for $\mathcal{A} \vdash \mathcal{B}$ and one for $\mathcal{B} \vdash \mathcal{A}$	give an interpretation in which \mathcal{A} and \mathcal{B} have different truth values
Are $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ jointly consistent ?	give an interpretation in which all of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ are true	prove a contradiction from assumptions $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$
Is $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \therefore \mathcal{C}$ valid ?	give a proof with assumptions $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ and concluding with \mathcal{C}	give an interpretation in which each of $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n$ is true and \mathcal{C} is false

Appendices

A Alternative terminology and notation

1.1 Alternative terminology

Sentential logic. The study of Sentential goes by other names. The name we have given it – sentential logic – derives from the fact that it deals with whole sentences as its most basic building blocks. Other features motivate different names. Sometimes it is called *truth-functional logic*, because it deals only with assignments of truth and falsity to sentences, and its connectives are all truth-functional. Sometimes it is called *propositional logic*, which strikes me as a misleading choice. This may sometimes be innocent, as some people use ‘proposition’ to mean ‘sentence’. However, noting that different sentences can mean the same thing, many people use the term ‘proposition’ as a useful way of referring to the meaning of a sentence, what that sentence expresses. In *this* sense of ‘proposition’, ‘propositional logic’ is not a good name for the study of Sentential, since synonymous but not truth-functionally equivalent sentences like ‘Vixens are bold’ and ‘Female foxes are bold’ will be logically distinguished even though they express the same proposition.

Quantifier logic. The study of Quantifier goes by other names. Sometimes it is called *predicate logic*, because it allows us to apply predicates to objects. Sometimes it is called *first-order logic*, because it makes use only of quantifiers over objects, and variables that can be substituted for constants. This is to be distinguished from *higher-order logic*, which introduces quantification over properties, and variables that can be substituted for predicates. (This device would allow us to formalise such sentences as ‘Jane and Kane are alike in *every respect*’, treating the italicised phrase as a quantifier over ‘respects’, i.e., properties. This results in something like $\forall P(Pj \leftrightarrow Pk)$, which is *not* a sentence of Quantifier.)

Atomic Sentences. Some texts call atomic sentences *sentence letters*.

Formulas. Some texts call formulas *well-formed formulas*. Since ‘well-formed formula’ is such a long and cumbersome phrase, they then abbreviate this as *wff*. This is both barbarous and unnecessary (such texts do not make any important use of the contrasting class of ‘ill-formed formulas’). I have stuck with ‘formula’.

In §6, I defined *sentences* of Sentential. These are also sometimes called ‘formulas’ (or ‘well-formed formulas’) since in Sentential, unlike Quantifier, there is no distinction between a formula and a sentence.

Valuations. Some texts call valuations *truth-assignments*; others call them *structures*.

Expressive adequacy. Some texts describe Sentential as *truth-functionally complete*, rather than expressively adequate.

n-place predicates. I have called predicates ‘one-place’, ‘two-place’, ‘three-place’, etc. Other texts respectively call them ‘monadic’, ‘dyadic’, ‘triadic’, etc. Still other texts call them ‘unary’, ‘binary’, ‘ternary’, etc.

Names. In Quantifier, I have used ‘*a*’, ‘*b*’, ‘*c*’, for names. Some texts call these ‘constants’, because they have a constant referent in a given interpretation, as opposed to variables which have variable referents. Other texts do not mark any difference between names and variables in the syntax. Those texts focus simply on whether the symbol occurs *bound* or *unbound*.

Domains. Some texts describe a domain as a ‘domain of discourse’, or a ‘universe of discourse’.

Interpretations. Some texts call interpretations *models*; others call them *structures*.

1.2 Alternative 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.

This appendix presents some common symbols, so that you can recognise them if you encounter them in an article or in another book.

Negation. Two commonly used symbols are the *not sign*, ‘ \neg ’, and the *tilde operator*, ‘ \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 ’.

Some texts use an overline to indicate negation, so that ‘ $\overline{\mathcal{A}}$ ’ expresses the same thing as ‘ $\neg\mathcal{A}$ ’. This is clear enough if \mathcal{A} is an atomic sentence, but quickly becomes cumbersome if we attempt to nest negations: ‘ $\neg(A \wedge \neg(\neg\neg B \wedge C))$ ’ becomes the unwieldy

$$\overline{A \wedge (\overline{\overline{B \wedge C}})}$$

Disjunction. The symbol ‘ \vee ’ is typically used to symbolize inclusive disjunction. One etymology is from the Latin word ‘vel’, meaning ‘or’.

Conjunction. Conjunction is often symbolized with the *ampersand*, ‘&’. The ampersand is a decorative form of the Latin word ‘et’, which means ‘and’. (Its etymology still lingers in certain fonts, particularly in italic fonts; thus an italic ampersand might appear as ‘&’.) Using this symbol is not recommended, since it is commonly used in natural English writing (e.g., ‘Smith & Sons’). As a symbol in a formal system, the ampersand is not the English word ‘&’, so it is much neater to use a completely different symbol. The most common choice now is ‘ \wedge ’, which 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 ’. (These are often texts that use the overlining notation for negation. Such texts often involve languages in which conjunction and negation are the only connectives, and they typically also dispense with parentheses which are unnecessary in such austere languages.)

Material Conditional. There are two common symbols for the material conditional: the *arrow*, ‘ \rightarrow ’, and the *hook*, ‘ \supset ’. Rarely you might see ‘ \Rightarrow ’.

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 symbolised as a rotated ‘ A ’, and the existential quantifier as a rotated ‘ E ’. In some texts, there is no separate symbol for the universal quantifier. Instead, the variable is just written in parentheses in front of the formula that it binds. For example, they might write ‘ $(x)Px$ ’ where we would write ‘ $\forall xPx$ ’.

These alternative notations are summarised below:

negation	\neg, \sim, \overline{A}
conjunction	$\wedge, \&, \cdot$
disjunction	\vee
conditional	$\rightarrow, \supset, \Rightarrow$
biconditional	\leftrightarrow, \equiv
universal quantifier	$\forall x, (x)$

B Quick reference

2.1 Schematic Truth Tables

\mathcal{A}	$\neg\mathcal{A}$	\mathcal{A}	\mathcal{B}	$\mathcal{A} \wedge \mathcal{B}$	$\mathcal{A} \vee \mathcal{B}$	$\mathcal{A} \rightarrow \mathcal{B}$	$\mathcal{A} \leftrightarrow \mathcal{B}$
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

2.2 Symbolisation

Sentential Connectives

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 $(\neg P \wedge \neg Q)$
Both P, and Q	$(P \wedge Q)$
If P, then Q	$(P \rightarrow Q)$
P only if Q	$(P \rightarrow Q)$
P if and only if Q	$(P \leftrightarrow Q)$
P unless Q	$(P \vee Q)$

Predicates

All Fs are Gs	$\forall x(Fx \rightarrow Gx)$
Some Fs are Gs	$\exists x(Fx \wedge Gx)$
Not all Fs are Gs	$\neg\forall x(Fx \rightarrow Gx)$ or $\exists x(Fx \wedge \neg Gx)$
No Fs are Gs	$\forall x(Fx \rightarrow \neg Gx)$ or $\neg\exists x(Fx \wedge Gx)$

Identity

Only c is G	$\forall x(Gx \leftrightarrow x = c)$
Everything besides c is G	$\forall x(\neg x = c \rightarrow Gx)$
The F is G	$\exists x(Fx \wedge \forall y(Fy \rightarrow x = y) \wedge Gx)$
It is not the case that the F is G	$\neg \exists x(Fx \wedge \forall y(Fy \rightarrow x = y) \wedge Gx)$
The F is non- G	$\exists x(Fx \wedge \forall y(Fy \rightarrow x = y) \wedge \neg Gx)$

2.3 Using identity to symbolize quantities

There are at least _____ Fs.

- one: $\exists xFx$
 two: $\exists x_1 \exists x_2 (Fx_1 \wedge Fx_2 \wedge \neg x_1 = x_2)$
 three: $\exists x_1 \exists x_2 \exists x_3 (Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3)$
 n : $\exists x_1 \dots \exists x_n (Fx_1 \wedge \dots \wedge Fx_n \wedge \neg x_1 = x_2 \wedge \dots \wedge \neg x_{n-1} = x_n)$

There are at most _____ Fs.

One way to say ‘there are at most n Fs’ is to put a negation sign in front of the symbolisation for ‘there are at least $n + 1$ Fs’. Equivalently, we can offer:

- one: $\forall x_1 \forall x_2 ((Fx_1 \wedge Fx_2) \rightarrow x_1 = x_2)$
 two: $\forall x_1 \forall x_2 \forall x_3 ((Fx_1 \wedge Fx_2 \wedge Fx_3) \rightarrow (x_1 = x_2 \vee x_1 = x_3 \vee x_2 = x_3))$
 three: $\forall x_1 \forall x_2 \forall x_3 \forall x_4 ((Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge Fx_4) \rightarrow$
 $(x_1 = x_2 \vee x_1 = x_3 \vee x_1 = x_4 \vee x_2 = x_3 \vee x_2 = x_4 \vee x_3 = x_4))$
 n : $\forall x_1 \dots \forall x_{n+1} ((Fx_1 \wedge \dots \wedge Fx_{n+1}) \rightarrow (x_1 = x_2 \vee \dots \vee x_n = x_{n+1}))$

There are exactly _____ Fs.

One way to say ‘there are exactly n Fs’ is to conjoin two of the symbolizations above and say ‘there are at least n Fs and there are at most n Fs.’ The following equivalent formulae are shorter:

- zero: $\forall x \neg Fx$
 one: $\exists x (Fx \wedge \forall y (Fy \rightarrow x = y))$
 two: $\exists x_1 \exists x_2 (Fx_1 \wedge Fx_2 \wedge \neg x_1 = x_2 \wedge \forall y (Fy \rightarrow (y = x_1 \vee y = x_2)))$
 three: $\exists x_1 \exists x_2 \exists x_3 (Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3 \wedge$
 $\forall y (Fy \rightarrow (y = x_1 \vee y = x_2 \vee y = x_3)))$
 n : $\exists x_1 \dots \exists x_n (Fx_1 \wedge \dots \wedge Fx_n \wedge \neg x_1 = x_2 \wedge \dots \wedge \neg x_{n-1} = x_n \wedge$
 $\forall y (Fy \rightarrow (y = x_1 \vee \dots \vee y = x_n)))$

2.4 Basic deduction rules for Sentential

Conjunction

$$\begin{array}{l|l} m & \mathcal{A} \\ n & \mathcal{B} \\ & \hline & \mathcal{A} \wedge \mathcal{B} \quad \wedge I\ m, n \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \wedge \mathcal{B} \\ & \hline & \mathcal{A} \quad \wedge E\ m \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \wedge \mathcal{B} \\ & \hline & \mathcal{B} \quad \wedge E\ m \end{array}$$

Conditional

$$\begin{array}{l|l|l} i & & \mathcal{A} \\ j & & \mathcal{B} \\ & & \hline & \mathcal{A} \rightarrow \mathcal{B} \quad \rightarrow I\ i-j \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \rightarrow \mathcal{B} \\ n & \mathcal{A} \\ & \hline & \mathcal{B} \quad \rightarrow E\ m, n \end{array}$$

Negation

$$\begin{array}{l|l|l} i & & \mathcal{A} \\ j & & \mathcal{B} \\ k & & \neg \mathcal{B} \\ & & \hline & \neg \mathcal{A} \quad \neg I\ i-j, i-k \end{array}$$

$$\begin{array}{l|l} m & \neg \neg \mathcal{A} \\ & \hline & \mathcal{A} \quad \neg \neg E\ m \end{array}$$

Disjunction

$$\begin{array}{l|l} m & \mathcal{A} \\ & \hline & \mathcal{A} \vee \mathcal{B} \quad \vee I\ m \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \\ & \hline & \mathcal{B} \vee \mathcal{A} \quad \vee I\ m \end{array}$$

$$\begin{array}{l|l|l} m & \mathcal{A} \vee \mathcal{B} \\ i & & \mathcal{A} \\ j & & \mathcal{C} \\ k & & \mathcal{B} \\ l & & \mathcal{C} \\ & & \hline & \mathcal{C} \quad \vee E\ m, i-j, k-l \end{array}$$

Biconditional

$$\begin{array}{l|l|l} i & & \mathcal{A} \\ j & & \mathcal{B} \\ k & & \mathcal{B} \\ l & & \mathcal{A} \\ & & \hline & \mathcal{A} \leftrightarrow \mathcal{B} \quad \leftrightarrow I\ i-j, k-l \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \leftrightarrow \mathcal{B} \\ n & \mathcal{A} \\ & \hline & \mathcal{B} \quad \leftrightarrow E\ m, n \end{array}$$

$$\begin{array}{l|l} m & \mathcal{A} \leftrightarrow \mathcal{B} \\ n & \mathcal{B} \\ & \hline & \mathcal{A} \quad \leftrightarrow E\ m, n \end{array}$$

2.5 Derived rules for Sentential

Reiteration

$$\begin{array}{l|l}
 m & \mathcal{A} \\
 & \mathcal{A} \quad Rm
 \end{array}$$

Disjunctive syllogism

$$\begin{array}{l|l}
 m & \mathcal{A} \vee \mathcal{B} \\
 n & \neg \mathcal{A} \\
 & \mathcal{B} \quad DS\ m, n
 \end{array}$$

$$\begin{array}{l|l}
 m & \mathcal{A} \vee \mathcal{B} \\
 n & \neg \mathcal{B} \\
 & \mathcal{A} \quad DS\ m, n
 \end{array}$$

Modus Tollens

$$\begin{array}{l|l}
 m & \mathcal{A} \rightarrow \mathcal{B} \\
 n & \neg \mathcal{B} \\
 & \neg \mathcal{A} \quad MT\ m, n
 \end{array}$$

Tertium non datur

$$\begin{array}{l|l}
 i & \mathcal{A} \\
 j & \mathcal{B} \\
 k & \neg \mathcal{A} \\
 l & \mathcal{B} \\
 & \mathcal{B} \quad TND\ i-j, k-l
 \end{array}$$

De Morgan Rules

$$\begin{array}{l|l}
 m & \neg(\mathcal{A} \vee \mathcal{B}) \\
 & \neg \mathcal{A} \wedge \neg \mathcal{B} \quad DeM\ m
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg \mathcal{A} \wedge \neg \mathcal{B} \\
 & \neg(\mathcal{A} \vee \mathcal{B}) \quad DeM\ m
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg(\mathcal{A} \wedge \mathcal{B}) \\
 & \neg \mathcal{A} \vee \neg \mathcal{B} \quad DeM\ m
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg \mathcal{A} \vee \neg \mathcal{B} \\
 & \neg(\mathcal{A} \wedge \mathcal{B}) \quad DeM\ m
 \end{array}$$

2.6 Basic deduction rules for Quantifier

Universal elimination

$$m \left| \begin{array}{l} \forall x \mathcal{A} \\ \mathcal{A}|_{c \sim x} \end{array} \right. \quad \forall E \ m$$

Existential introduction

$$m \left| \begin{array}{l} \mathcal{A}|_{c \sim x} \\ \exists x \mathcal{A} \end{array} \right. \quad \exists I \ m$$

Universal introduction

$$m \left| \begin{array}{l} \mathcal{A}|_{c \sim x} \\ \forall x \mathcal{A} \end{array} \right. \quad \forall I \ m$$

c must not occur in any undischarged assumption, or in \mathcal{A}

Existential elimination

$$m \left| \begin{array}{l} \exists x \mathcal{A} \\ i \left| \begin{array}{l} \mathcal{A}|_{c \sim x} \\ \hline \mathcal{B} \end{array} \right. \\ j \left| \begin{array}{l} \mathcal{B} \end{array} \right. \end{array} \right. \quad \exists E \ m, i-j$$

c must not occur in any undischarged assumption, in $\exists x \mathcal{A}$, or in \mathcal{B}

Identity introduction

$$\left| \begin{array}{l} c = c \end{array} \right. \quad =I$$

Identity elimination

$$m \left| \begin{array}{l} a = b \\ n \left| \begin{array}{l} \mathcal{A}|_{a \sim x} \\ \mathcal{A}|_{b \sim x} \end{array} \right. \end{array} \right. \quad =E \ m, n$$

2.7 Derived rules for Quantifier

$$m \left| \begin{array}{l} \forall x \neg \mathcal{A} \\ \neg \exists x \mathcal{A} \end{array} \right. \quad \text{CQ 1}$$

$$m \left| \begin{array}{l} \neg \exists x \mathcal{A} \\ \forall x \neg \mathcal{A} \end{array} \right. \quad \text{CQ 1}$$

$$m \left| \begin{array}{l} \exists x \neg \mathcal{A} \\ \neg \forall x \mathcal{A} \end{array} \right. \quad \text{CQ 1}$$

$$m \left| \begin{array}{l} \neg \forall x \mathcal{A} \\ \exists x \neg \mathcal{A} \end{array} \right. \quad \text{CQ 1}$$

$$m \left| \begin{array}{l} a = b \\ n \left| \begin{array}{l} \mathcal{A}|_{b \sim x} \\ \mathcal{A}|_{a \sim x} \end{array} \right. \end{array} \right. \quad =ES \ m, n$$

Acknowledgements

P.D. Magnus would like to thank the people who made this project possible. Notable among these are Cristyn Magnus, who read many early drafts; Aaron Schiller, who was an early adopter and provided considerable, helpful feedback; and Bin Kang, Craig Erb, Nathan Carter, Wes McMichael, Selva Samuel, Dave Krueger, Brandon Lee, and the students of Introduction to Logic, who detected various errors in previous versions of the book.

Tim Button would like to thank P.D. Magnus for his extraordinary act of generosity, in making forallx available to everyone. Thanks also to Alfredo Manfredini Böhm, Sam Brain, Felicity Davies, Emily Dyson, Phoebe Hill, Richard Jennings, Justin Niven, and Igor Stojanovic for noticing errata in earlier versions.

Antony Eagle would like to thank P.D. Magnus and Tim Button for their work from which this text derives, and Jon Opie for comments.

About the Authors

P.D. Magnus is a professor of philosophy in Albany, New York. His primary research is in the philosophy of science. www.fecundity.com/job/

Tim Button is a University Lecturer, and Fellow of St John's College, at the University of Cambridge. His first book, *The Limits of Realism*, was published by Oxford University Press in 2013. nottub.com

Antony Eagle is Senior Lecturer in Philosophy at the University of Adelaide. His research interests include metaphysics, philosophy of probability, philosophy of physics, and philosophy of language. antonyeagle.org

In the Introduction to his *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 readers are forgiven if they take a break for snacks after *two* readings.