

forall $\lambda$ :Cambridge

P.D. Magnus  
*University at Albany, State University of New York*

Modified for the Cambridge Course by:  
Tim Button  
*University of Cambridge*

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, Igor Stojanovic and Rob Trueman for noticing mistakes in earlier versions; and to Rob, again, for discussions on how to label gaps hygienically (see §16).

© 2005–2018 by P.D. Magnus and Tim Button. Some rights reserved.

This book is based upon P.D. Magnus’s `forallx`, available at [fecundity.com/logic](https://fecundity.com/logic), which was released under a Creative Commons license (CC BY 4.0; full details are available at [creativecommons.org/licenses/by/4.0/](https://creativecommons.org/licenses/by/4.0/)). The following is a human-readable summary of (and not a substitute for) that license:

**License details.** You are free to:

- *Share*: copy and redistribute the material in any medium or format
- *Adapt*: remix, transform, and build upon the material for any purpose, even commercially.

Under the following terms:

- *Attribution*: You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

In accordance with this license, Tim Button has made changes to P.D. Magnus’s original text, and added new material, and he offers `forallx:Cambridge` under the same CC BY 4.0 license. This copy of `forallx:Cambridge` was released on Saturday, 22 September 2018. The most recent version is available at [nottub.com](https://nottub.com).

Typeset in X<sub>Y</sub>-L<sup>A</sup>T<sub>E</sub>X, using Libertinus Serif and Libertinus Sans. The style for typesetting proofs is based on `fitch.sty` (v0.4) by Peter Selinger, University of Ottawa.

# Contents

<b>1</b>	<b>Key notions</b>	<b>1</b>
1	Arguments	2
2	Valid arguments	4
3	Other logical notions	7
<b>2</b>	<b>Truth-functional logic</b>	<b>10</b>
4	First steps to symbolisation	11
5	Connectives	14
6	Sentences of TFL	23
7	Use and mention	27
<b>3</b>	<b>Truth tables</b>	<b>31</b>
8	Characteristic truth tables	32
9	Truth-functional connectives	34
10	Complete truth tables	37
11	Semantic concepts	42
12	Truth table shortcuts	47
13	Partial truth tables	50
<b>4</b>	<b>First-order logic</b>	<b>53</b>
14	Building blocks of FOL	54
15	Sentences with one quantifier	58
16	Multiple generality	66
17	Identity	73
18	Definite descriptions	77
19	Sentences of FOL	82
<b>5</b>	<b>Interpretations</b>	<b>87</b>
20	Extensionality	88
21	Truth in FOL	92
22	Semantic concepts	97
23	Using interpretations	98
24	Reasoning about all interpretations	103

---

<b>6</b>	<b>Natural deduction for TFL</b>	<b>106</b>
25	The very idea of natural deduction	107
26	Basic rules for TFL	109
27	Additional rules for TFL	125
28	Proof-theoretic concepts	129
29	Proof strategies	132
30	Derived rules	133
<b>7</b>	<b>Natural deduction for FOL</b>	<b>137</b>
31	Basic rules for FOL	138
32	Conversion of quantifiers	147
33	Rules for identity	149
34	Derived rules	152
35	Proof-theoretic concepts and semantic concepts	153
	<b>Appendices</b>	<b>157</b>
36	Symbolic notation	157
37	Alternative proof systems	159
38	Quick reference	163

**1**

# **Key notions**

# Arguments

# 1

Logic is the business of evaluating arguments; sorting the good from the bad.

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 just disagreements.

An argument, as we shall understand it, is something more like this:

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 you believe the premises, then the argument (perhaps) provides you with a reason to believe the conclusion.

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.

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.

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 guide, these words are often used to indicate an argument’s conclusion:

so, therefore, hence, thus, accordingly, consequently

By contrast, 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 exercises that review and explore the material covered in the chapter. There is no substitute for actually working through some problems, because learning logic is more about developing a way of thinking than it is about memorising facts.

So here's the first exercise. 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.

# Valid arguments

## 2

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 the *Cathedra Romana*.

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. The Pope's throne is only ever occupied by a hat-wearing man. Second: the conclusion does not follow from the premise of the argument. Even if there were a bassoon-playing dragon in the Pope's throne, 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.

The general point is that, 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.

It is often very important to determine whether or not the premises of an argument are true. But that is normally a task best left to experts in the relevant 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.

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.)

## 2.2 Validity

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 **VALID** if and only if it is impossible for all of the premises to be true and the conclusion false.

The crucial thing about a valid argument is that, if the premises are true then (necessarily) the conclusion is too. 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 are true, *then* the conclusion just has to be true. So the argument is valid.

This highlights that valid arguments do not need to have true premises or 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. So the argument is invalid.

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

## 2.3 Inductive arguments

Many good arguments are 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.

In January 2001, it rained in London.

In January 2002, 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 2003, it rained in London; In January 2004.... But, however many premises of this form we add, the argument will remain invalid. 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?

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. 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.

# Other logical notions

# 3

In §2, I introduced the idea of a valid argument. This is one of the most important ideas in logic. In this section, I will introduce are some similarly important ideas.

## 3.1 Truth values

As I said in §1, arguments consist of premises and a conclusion. But many kinds of English sentence *cannot* be used as 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 are not *assertoric*: 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 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!)

### 3.3 Necessity and contingency

In assessing arguments for validity, 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. Either it is raining here, or it is not.
3. 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 is true. Regardless of what the weather is like, it is either raining or it is not. That is a NECESSARY TRUTH.

Equally, you do not need to check the weather to determine whether or not sentence 3 is true. It must be false, simply as a matter of logic. It might be raining here and not raining across town; it might be raining now but stop raining even as you 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.

Something which is capable of being true or false, but which is neither necessarily true nor necessarily false, is CONTINGENT.

### 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 G1–G4 in this section (about giraffes, gorillas and martians in the wild animal park), and consider each of the following:

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

Which are jointly consistent? Which are jointly inconsistent?

C. Could there be:

1. A valid argument, the conclusion of which is necessarily false?
2. An invalid 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.

2

## Truth-functional logic

# First steps to symbolisation

# 4

## 4.1 Validity in virtue of form

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 looks like an excellent argument *structure*. Indeed, 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-(A and B)  
 A  
 So: not-B

The examples illustrate an important idea, which we might describe as *validity in virtue of form*. 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, I will develop a formal language which allows us to symbolise many arguments in such a way as to show that they are valid in virtue of their form. That language will be *truth-functional logic*, or TFL.

## 4.2 Validity for special reasons

There are plenty of arguments that are valid, but not for reasons relating to their form. 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 valid. But the validity is not related to the form of the argument. Here is an invalid argument with the same form:

Juanita is a vixen  
 So: Juanita is a cathedral

This might suggest that the validity 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 valid. Equally, consider the argument:

The sculpture is green all over.  
 So: The sculpture is not red all over.

Again, it seems impossible for the premise to be true and the conclusion false, for nothing can be both green all over and red all over. So the argument is valid. But here is an invalid argument with the same form:

The sculpture is green all over.  
 So: The sculpture is not shiny all over.

The argument is invalid, 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 validity of the first 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 valid.

The important moral can be stated as follows. *At best, TFL will help us to understand arguments that are valid due to their form.*

### 4.3 Atomic sentences

I started isolating the form of an argument, in §4.1, by replacing *subsences* of sentences with individual 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, TFL, 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 uppercase italic letters for atomic sentences of TFL. 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 TFL:

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

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 think of the atomic sentence of TFL, ‘A’, as symbolising the English sentence ‘It is raining outside’, and the atomic sentence of TFL, ‘C’, as symbolising 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 TFL. From the point of view of TFL, an atomic sentence is just a letter. It can be used to build more complex sentences, but it cannot be taken apart.

# Connectives

# 5

In the previous section, we considered symbolising fairly basic English sentences with atomic sentences of TFL. 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 TFL, we shall make use of logical connectives to build complex sentences from atomic components. There are five logical connectives in TFL. 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 ‘ $\neg B$ ’.

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

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 TFL 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 TFL-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 TFL.

## 5.2 Conjunction

Consider these sentences:

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

We will need separate atomic sentences of TFL 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 TFL.

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 TFL. 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 TFL 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 TFL 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’. TFL does not (and cannot) preserve these nuances.

Sentence 15 raises similar issues. There is a contrastive structure, but this is not something that TFL 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 TFL 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 have put *brackets* around the conjunctions. My reasons can be brought out by thinking about negation interacts 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 keep track of things like the *scope* of the negation.

### 5.3 Disjunction

Consider these sentences:

18. Either Denison will play golf with me, or he will watch movies.
19. 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 18 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 19 is only slightly more complicated. There are two subjects, but the English sentence only gives the verb once. However, we can paraphrase sentence 19 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  $(A \vee B)$  if it can be paraphrased in English as ‘Either..., or...’

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 19, above. I might play golf with Denison, with Ellery, or with both Denison and Ellery. Sentence 19 merely says that I will play with *at least* one of them. This is an INCLUSIVE OR. The TFL symbol ‘ $\vee$ ’ always symbolises an *inclusive or*.

It will also help to see how negation interacts with disjunction. Consider:

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

Using the same symbolisation key as before, sentence 20 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 TFL, 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 20 itself is symbolised by ' $(\neg S_1 \vee \neg S_2)$ '.

Sentence 21 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 21 with ' $\neg(S_1 \vee S_2)$ '.

Sentence 22 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 that both you get soup and 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 ' $\wedge$ '. So sentence 22 can be symbolised as ' $((S_1 \vee S_2) \wedge \neg(S_1 \wedge S_2))$ '.

This last example shows something important. Although the TFL symbol ' $\vee$ ' always symbolises *inclusive or*, we can symbolise an *exclusive or* in TFL. We just have to use a few other symbols too.

## 5.4 Conditional

Consider these sentences:

23. If Jean is in Paris, then Jean is in France.
24. 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 23 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 23 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 24 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 23. That would be a mistake. My knowledge of geography tells me that sentence 23 is unproblematically true: there is no way for Jean to be in Paris that doesn't involve Jean being in France. But sentence 24 is not so straightforward: were Jean in Dieppe, Lyons, or Toulouse, Jean would be in France without being in Paris, thereby rendering sentence 24 false. Since geography alone dictates the truth of sentence 23, whereas travel plans (say) are needed to know the truth of sentence 24, they must mean different things.

In fact, sentence 24 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  $(A \rightarrow B)$  if it can be paraphrased in English as 'If A, then B' or 'A only if B'.

In fact, the conditional can represent many English expressions. Consider:

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

28. It is a sufficient condition on Jean's being in France that she be in Paris.

If we think about it, 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 lose a huge amount when we use ' $\rightarrow$ ' to symbolise English conditionals. I shall return to this in §§9.3 and 11.5.

## 5.5 Biconditional

Consider these sentences:

- 29. Shergar is a horse only if he is a mammal
- 30. Shergar is a horse if he is a mammal
- 31. 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

For reasons discussed above, sentence 29 can be symbolised by ' $(H \rightarrow M)$ '.

Sentence 30 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 31 says something stronger than either 29 or 30. 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 29 and 30. So we can symbolise it as ' $((H \rightarrow M) \wedge (M \rightarrow H))$ '. We call this a **BICONDITIONAL**, because it amounts to stating both directions of the conditional.

We could treat every biconditional this way. So, just as we do not need a new TFL symbol to deal with *exclusive or*, we do not really need a new TFL symbol to deal with biconditionals. However, we will use ' $\leftrightarrow$ ' to symbolise the biconditional. So we can symbolise sentence 31 with the TFL sentence ' $(H \leftrightarrow M)$ '.

The expression 'if and only if' occurs a lot in philosophy and logic. For brevity, it is often abbreviated with a single, snappy 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  $(A \leftrightarrow B)$  if it can be paraphrased in English as 'A iff B'; that is, as 'A if and only if B'.

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 greens, you won't get any pudding'. Suppose you ate your greens, but that your then parents refused to give you any pudding, on the grounds that they were only committed to the *conditional* (roughly 'if you get pudding, then you will have eaten your greens'), rather than the biconditional (roughly, 'you get pudding iff you eat your greens'). Well, a tantrum would rightly ensue. So, be aware of this when interpreting people; but in your own writing, make sure to use the biconditional iff you mean to.

## 5.6 Unless

We have now seen all of the connectives of TFL. We can use them together to symbolise many kinds of sentences. But some cases are harder than others. And a typically nasty case is the English-language connective ‘unless’:

32. Unless you wear a jacket, you will catch a cold.  
 33. You will catch a 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 a 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 TFL. For simplicity, then:

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 own writing. Unless you want to be ambiguous.

### Practice exercises

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

- $M$ : 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 TFL.

- 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 TFL.

- $E_1$ : 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, then he must be satisfied with his work.
11. It cannot be that Harrison is both an electrician and a firefighter.
12. Harrison and Ava are both firefighters if and only if neither of them is an electrician.

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

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 TFL.

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 TFL.

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?

# Sentences of TFL

# 6

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 TFL. 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, I will *define* exactly what will count as a sentence of TFL. This is one respect in which a formal language like TFL is more precise than a natural language like English.

## 6.1 Expressions

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

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

Define an EXPRESSION OF TFL as any string of symbols of TFL. So: write down any sequence of symbols of TFL, in any order, and you have an expression of TFL.

## 6.2 Sentences

Given what we just said, ‘ $(A \wedge B)$ ’ is an expression of TFL, and so is ‘ $\neg(\vee()) \wedge (\neg\neg())((B)$ ’. But the former is a *sentence*, and the latter is *gibberish*. We want some rules to tell us which TFL expressions are sentences.

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 of the sentences one by one.

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

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

1. Every atomic sentence is a sentence.
2. If  $A$  is a sentence, then  $\neg A$  is a sentence.
3. If  $A$  and  $B$  are sentences, then  $(A \wedge B)$  is a sentence.
4. If  $A$  and  $B$  are sentences, then  $(A \vee B)$  is a sentence.
5. If  $A$  and  $B$  are sentences, then  $(A \rightarrow B)$  is a sentence.
6. If  $A$  and  $B$  are sentences, then  $(A \leftrightarrow B)$  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, consider a recursive definition of the notion 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 TFL. 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 ultimately get down to atomic sentences, then we are ok.

Let's put this into practice by considering some examples.

Suppose we want to know whether or not  $\neg\neg\neg D$  is a sentence of TFL. 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 TFL, 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 *most recently*, when constructing the sentence. We call that the MAIN LOGICAL OPERATOR of the sentence. In the case of ‘ $\neg\neg\neg D$ ’, the main logical operator is the very first ‘ $\neg$ ’ sign. In the case of ‘ $(P \wedge \neg(\neg Q \vee R))$ ’, the main logical operator is ‘ $\wedge$ ’. In the case of ‘ $((\neg E \vee F) \rightarrow \neg\neg G)$ ’, the main logical operator is ‘ $\rightarrow$ ’.

As a general rule, you can find the main logical operator for a sentence by using the following method:

- If the first symbol in the sentence is ‘ $\neg$ ’, then that is the main logical operator
- Otherwise, start counting the brackets. For each open-bracket, i.e. ‘(’, add 1; for each closing-bracket, i.e. ‘)’, subtract 1. When your count is at exactly 1, the first operator you hit (*apart* from a ‘ $\neg$ ’) is the main logical operator.

(Note: if you do use this method, then make sure to included *all* the brackets in the sentence, rather than omitting some as per the conventions of §6.3!)

The recursive structure of sentences in TFL 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. I will return to this in chapter 3.

The recursive structure of sentences in TFL 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 logical operator. For example, consider this sentence:

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

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 logical operator. So the scope of the negation is just ‘ $\neg(R \wedge B)$ ’. More generally:

The SCOPE of a connective in a sentence is the subsentence for which that connective is the main logical operator.

### 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 TFL, 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 another 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 square

brackets, '[' and ']', instead of rounded ones. So there is no logical difference between ' $(P \vee Q)$ ' and ' $[P \vee Q]$ ', for example.

Combining these two conventions, we can rewrite the unwieldy sentence

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

rather more simply 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 TFL, strictly speaking? (b) Is it a sentence of TFL, allowing for our relaxed bracketing conventions?

1.  $(A)$
2.  $\mathcal{J}_{374} \vee \neg \mathcal{J}_{374}$
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 TFL that contain no atomic sentences? Explain your answer.

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

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

D. In §6.2, I provided a method for finding the main logical connective in any sentence. Try to explain *why* the method always works.

# Use and mention

# 7

In this chapter, I have talked a lot *about* sentences. I now want to pause, to explain an important—and very general—point about talking *about* languages.

## 7.1 Quotation conventions

Consider these two sentences:

- Elizabeth is the Queen.
- The expression ‘Elizabeth’ is composed of one uppercase letter and eight lowercase letters

When we want to talk about the Queen, we *use* her name. When we want to talk about the Queen’s name, we *mention* that name, which we do 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. 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:

- ‘Elizabeth’ is the Queen.

says that some *expression* is the Queen. And that’s false. The *woman* is the Queen; her *name* isn’t. Conversely, this sentence:

- Elizabeth is composed of one uppercase letter and eight lowercase letters.

also says something false: Elizabeth is a woman, made of meat rather than letters. One final example:

- “‘Elizabeth’” names ‘Elizabeth’.

On the left-hand-side, 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 reported speech (which I might use in talking about Descartes’s ‘cogito ergo sum’). They indicate that you are moving from talking about an object, to talking about a name of that object.

## 7.2 Object language and metalanguage

These general quotation conventions are very important for us. After all, we are describing a formal language here, TFL, and so we must often *mention* expressions from TFL.

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: TFL. The metalanguage is English. Not conversational English, perhaps, but English supplemented with some additional vocabulary to help us get along.

Now, I have used italic uppercase letters for atomic sentences of TFL:

$$A, B, C, Z, A_1, B_4, A_{25}, \mathcal{J}_{375}, \dots$$

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

- *D* is an atomic sentence of TFL.

Obviously, I am trying to come out with an English sentence that says something about the object language (TFL). But '*D*' is a sentence of TFL, 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 TFL.

The general point is that, whenever we want to talk in English about some specific expression of TFL, 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 Fonts, quotation marks, and concatenation

However, we do not just want to talk about *specific* expressions of TFL. We also want to be able to talk about *any arbitrary* sentence of TFL. Indeed, I had to do this in §6, when I presented the recursive definition of a sentence of TFL. I used bold-italic-fonts to do this:

$$\mathbf{A, B, C, D, \dots}$$

These symbols do not belong to TFL. Rather, they are part of our (augmented) metalanguage that we use to talk about *any* expression of TFL. To explain why we need them, recall the second clause of the recursive definition of a sentence of TFL:

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:

‘ $A$ ’ is a symbol in augmented English, and we use it to talk about expressions of TFL. ‘ $A$ ’ is a particular atomic sentence of TFL.

But all of this raises a further complication, concerning quotation conventions. I did not include any quotation marks in the third clause of our recursive definition. Should I have done so?

The problem is that the expression on the right-hand-side of the rule, i.e. ‘ $\neg A$ ’ is not a sentence of English, since it contains ‘ $\neg$ ’. So we might try to write:

- 3'. If  $A$  is a sentence, then ‘ $\neg A$ ’ is a sentence.

But this is no good: ‘ $\neg A$ ’ is not a TFL sentence, since ‘ $A$ ’ is a symbol of (augmented) English rather than a symbol of TFL.

What we really want to say is something like this:

- 3''. If  $A$  is a sentence, then the result of concatenating the symbol ‘ $\neg$ ’ with the sentence  $A$  is itself 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 ‘ $\neg A$ ’ should simply be read *directly* in terms of rules for concatenation. So, *officially*, the metalanguage expression ‘ $\neg A$ ’ simply abbreviates:

the result of concatenating the symbol ‘ $\neg$ ’ with the sentence  $A$

and similarly, for expressions like ‘ $(A \wedge B)$ ’, ‘ $(A \vee B)$ ’, etc.

## 7.4 Quotation conventions for arguments

One of our main purposes for using TFL 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 TFL.

Or rather, we can use TFL to symbolise each of the *sentences* used in an English argument. However, TFL itself has no way to flag some of them as the *premises* and another as the *conclusion* of an argument. (Contrast this with natural English, which uses words like ‘so’, ‘therefore’, etc., to mark that a sentence is the *conclusion* of an argument.)

So, we need another bit of notation. Suppose we want to symbolise the premises of an argument with  $A_1, \dots, A_n$  and the conclusion with  $C$ . Then we will write:

$$A_1, \dots, A_n \therefore C$$

The role of the symbol ‘ $\therefore$ ’ is simply to indicate which sentences are the premises and which is the conclusion.

Strictly, the symbol ‘ $\therefore$ ’ will not be a part of the object language, but of the *metalanguage*. As such, one might think that we would need to put quote-marks around the TFL-sentences which flank it. That is a sensible thought, but adding these quote-marks would make things harder to read. Moreover—and as above—recall that *we* are stipulating some new conventions. So, we can simply stipulate that these quote-marks are unnecessary. That is, we can simply write:

$$A, A \rightarrow B \therefore B$$

*without any quotation marks*, to indicate an argument whose premises are (symbolised by) ‘ $A$ ’ and ‘ $A \rightarrow B$ ’ and whose conclusion is (symbolised by) ‘ $B$ ’.

3

## Truth tables

# Characteristic truth tables

# 8

Any non-atomic sentence of TFL is composed of atomic sentences with sentential connectives. The truth value of the compound sentence depends only on the truth value of the atomic sentences that comprise it. In order to know the truth value of ' $(D \wedge E)$ ', for instance, you only need to know the truth value of ' $D$ ' and the truth value of ' $E$ '.

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

**Negation.** For any sentence  $A$ : If  $A$  is true, then  $\neg A$  is false; and if  $\neg A$  is true, then  $A$  is false. We can summarize this in the *characteristic truth table* for negation:

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

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

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

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

**Disjunction.** Recall that ' $\vee$ ' always represents inclusive or. So, for any sentences  $A$  and  $B$ ,  $A \vee B$  is true if and only if either  $A$  or  $B$  is true. We can summarize this in the characteristic truth table for disjunction:

$A$	$B$	$A \vee B$
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 big old mess in TFL. Exactly how much of a mess they are is a matter of *philosophical* contention. I shall discuss a few of the subtleties in §§9.3 and 11.5. For now, I am going to stipulate the following:  $A \rightarrow B$  is false if and only if  $A$  is true and  $B$  is false. We can summarize this with a characteristic truth table for the conditional.

$A$	$B$	$A \rightarrow B$
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;  $A \rightarrow B$  and  $B \rightarrow A$  have different truth tables.

**Biconditional.** Since a biconditional is to be the same as the conjunction of a conditional running in both directions, the truth table for the biconditional should be:

$A$	$B$	$A \leftrightarrow B$
T	T	T
T	F	F
F	T	F
F	F	T

Unsurprisingly, the biconditional is symmetrical.

# Truth-functional connectives

9

## 9.1 The idea of truth-functionality

I want to introduce an important idea.

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

Every connective in TFL is truth-functional. 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 TFL sentence, we only need to know the truth value of its components.

This is what gives TFL its name: it is *truth-functional logic*.

Many languages use 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  $2 + 2 = 4$ , 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 not *truth-functional*.

## 9.2 Symbolising versus translating

All of the connectives of TFL 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 TFL, 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 TFL. As remarked in §5, TFL 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 TFL sentence, perhaps ' $F \wedge Q$ '.

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

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

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

Other textbooks will understand this as a stipulation that the TFL sentence ' $F$ ' should *mean* that Jon is fat, and that the TFL sentence ' $Q$ ' should *mean* that Jon is quick. But TFL just is totally unequipped to deal with *meaning*. The preceding symbolisation key is doing no more nor less than stipulating truth values for the TFL sentences ' $F$ ' and ' $Q$ '. We are laying down that ' $F$ ' should be true if Jon is fat (and false otherwise), and that ' $Q$ ' should be true if Jon is quick (and false otherwise.)

When we treat a TFL sentence as *symbolising* some English sentence, we are simply stipulating a truth value for that TFL sentence.

### 9.3 Indicative versus subjunctive conditionals

To bring home the point that TFL can *only* deal with truth functions, I want to say a little bit about conditionals. When I introduced the characteristic truth table for the material conditional in §8, I did not say anything to justify it. Let me now offer a justification, which follows Dorothy Edgington.<sup>1</sup>

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:



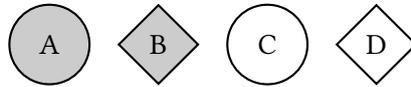
In this case, my claim is surely true. Shapes C and D are not grey, and so can hardly present *counterexamples* to my claim. Shape A *is* grey, but fortunately it is also circular. So my claim has no counterexamples. It must be true. And that means that each of the following *instances* of my claim must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

<sup>1</sup>Dorothy Edgington, 'Conditionals', 2014, in the *Stanford Encyclopedia of Philosophy* (<http://plato.stanford.edu/entries/conditionals/>).

- 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 this claim must also be false:

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

Now, recall that every connective of TFL 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 TFL can provide*. But is it any good, as a surrogate for the conditionals we use in everyday language? Consider two sentences:

1. If Hillary Clinton had won the 2016 US election, then she would have been the first female president of the US.
2. If Hillary Clinton had won the 2016 US election, then she 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. (Hillary did not win; she did not become the first female president of the US; and she did not fill with helium and float away.) So the truth value of the whole sentence is not uniquely determined by the truth value of the parts.

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—after all, Hillary Clinton lost the 2016 election—and then ask us to evaluate what *would* have happened in that case. Such considerations simply cannot be tackled using ‘ $\rightarrow$ ’.

I shall say more about the difficulties with conditionals in §11.5. For now, I shall content myself with the observation that ‘ $\rightarrow$ ’ is the only candidate for a truth-functional conditional, but that many English conditionals cannot be represented adequately using ‘ $\rightarrow$ ’. TFL is an intrinsically limited language. And you should not blithely assume that you can adequately symbolise an English ‘if..., then...’ with TFL’s ‘ $\rightarrow$ ’.

# Complete truth tables

So far, I have used symbolisation keys to assign truth values to TFL sentences *indirectly*. For example, we might say that the TFL sentence ‘*B*’ is to be true iff Big Ben is in London. Since Big Ben *is* in London, this symbolisation would make ‘*B*’ true. 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. Such stipulations are called *valuations*:

A VALUATION is any assignment of truth values to particular atomic sentences of TFL.

The power of truth tables lies in the following. Each row of a truth table represents a possible valuation. The complete 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. But all of this is easiest to explain by 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 valuations—which we can represent as follows:

<i>H</i>	<i>I</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:

<i>H</i>	<i>I</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,  $(A \wedge B)$ , with ‘*H*’ as *A* and with ‘*I*’ as *B*. The characteristic truth table for conjunction gives the truth conditions for *any* sentence of the form  $(A \wedge B)$ , whatever *A* and *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.

$H$	$I$	$A \wedge B$	$(H \wedge I) \rightarrow H$
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,  $A \rightarrow B$ , with ' $(H \wedge I)$ ' as  $A$  and with ' $H$ ' as  $B$ . 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:

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

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
T	F	T F F T T
F	T	F F T T F
F	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 logical operator* 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 characteristic truth table for negation. This is true even if the same letter is repeated many times, as in the sentence ‘ $[(C \leftrightarrow C) \rightarrow C] \wedge \neg(C \rightarrow C)$ ’. 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 F T F

Looking at the column underneath the main logical operator, 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.

There will be four lines in the complete truth table for a sentence containing two atomic sentences, as in the characteristic truth tables, or the truth table for ‘ $(H \wedge I) \rightarrow H$ ’.

There will be eight lines in the complete truth table for a sentence containing three atomic sentences, e.g.:

$M$	$N$	$P$	$M \wedge (N \vee P)$
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 sentence containing four different atomic sentences needs a truth table with 16 lines. Five atomic sentences, 32 lines. Six atomic sentences, 64 lines. And so on. To be perfectly general: a complete truth table with  $n$  atomic sentences must have  $2^n$  lines.

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.

### 10.3 More bracketing conventions

Consider these two sentences:

$$((A \wedge B) \wedge C)$$

$$(A \wedge (B \wedge C))$$

These have the same truth table. Consequently, it will never make any difference from the perspective of truth value—which is all that TFL 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.

# Semantic concepts

11

In the previous section, I introduced the idea of a valuation and showed how to determine the truth value of any TFL sentence, on any valuation, using a truth table. In this section, I 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 surrogates in TFL. We shall start with a surrogate for necessary truth.

A sentence is a TAUTOLOGY iff it is true on every valuation.

We can use truth tables to decide whether a sentence is a tautology. If the sentence is true on every line of its 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 TFL. One example is  $'2 + 2 = 4'$ . This *must* be true, but if we try to symbolise it in TFL, the best we can offer is an atomic sentence, and no atomic sentence is a tautology. Still, if we can adequately symbolise some English sentence using a TFL sentence which is a tautology, then that English sentence expresses a necessary truth.

We have a similar surrogate for necessary falsity:

A sentence is a CONTRADICTION iff it is false on every valuation.

We can use truth tables to decide whether a sentence is a contradiction. If the sentence is false on every line of its 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.

## 11.2 Tautological equivalence

Here is a similar, useful notion:

Sentences are TAUTOLOGICALLY 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 tautologically equivalent. Again, it is easy to test for tautological equivalence using truth tables. Consider the sentences  $\neg(P \vee Q)$  and  $\neg P \wedge \neg Q$ . Are they tautologically 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 logical operators; 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 tautologically 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:

Sentences are **JOINTLY TAUTOLOGICALLY CONSISTENT** iff there is some valuation which makes them all true.

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

### 11.4 Tautological entailment and validity

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

$A_1, A_2, \dots, A_n$  **TAUTOLOGICALLY ENTAIL**  $C$  iff no valuation of the relevant atomic sentences makes all of  $A_1, A_2, \dots, A_n$  true and  $C$  false.

Again, it is easy to test this with a truth table. To check whether ' $\neg L \rightarrow (J \vee L)$ ' and ' $\neg L$ ' tautologically 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	T	F	T
T	F	T	T	T
F	T	F	F	F
F	F	T	T	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 ' $\neg L$ ' tautologically entail ' $J$ '.

The notion of tautological entailment is deeply connected to the notion of validity:<sup>1</sup>

If  $A_1, \dots, A_n$  tautologically entail  $C$ , then  $A_1, \dots, A_n \therefore C$  is valid.

Here's why. If  $A_1, A_2, \dots, A_n$  tautologically entail  $C$ , then no valuation makes all of  $A_1, A_2, \dots, A_n$  true whilst making  $C$  false. So it is *logically impossible* for  $A_1, A_2, \dots, A_n$  all to be true whilst  $C$  is false. And this guarantees the validity of the argument with premises  $A_1, \dots, A_n$  and conclusion  $C$ .

In short, we have a way to test for the validity of English arguments. First, we symbolise them in TFL; then we test for tautological entailment using truth tables.

<sup>1</sup>Recall from §7 the use of ' $\therefore$ '; here. Without it, the information in this box could be written as follows. If  $A_1, \dots, A_n$  tautologically entail  $C$ , then the argument with premises  $A_1, \dots, A_n$  and conclusion  $C$  is valid.

## 11.5 The limits of these tests

This is 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 has four legs. So Daisy has more than two legs.

To symbolise this argument in TFL, we would have to use two different atomic sentences – perhaps ‘ $F$ ’ and ‘ $T$ ’ – for the premise and the conclusion respectively. Now, it is obvious that ‘ $F$ ’ does not tautologically entail ‘ $T$ ’. But the English argument is surely valid!

Second, consider the sentence:

2. Jan is neither bald nor not-bald.

To symbolise this sentence in TFL, 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 added ‘Jan is on the borderline of baldness!’

Third, consider the following sentence:

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

Symbolising this in TFL, we would offer something like ‘ $\neg(G \rightarrow M)$ ’. Now, ‘ $\neg(G \rightarrow M)$ ’ tautologically entails ‘ $G$ ’ (again, check this with a truth table). So if we symbolise sentence 3 in TFL, it seems to entail that God exists. But that’s strange: surely even the atheist can accept sentence 3, without contradicting herself!

In different ways, these three examples highlight some of the limits of working with a language (like TFL) 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 TFL, before we can adequately discuss its limits, and consider alternatives.

## 11.6 The double-turnstile

In what follows, we will use the notion of tautological entailment rather often. It will help us, then, to introduce a symbol that abbreviates it. Rather than saying that the TFL sentences  $A_1, A_2, \dots$  and  $A_n$  together tautologically entail  $C$ , we shall abbreviate this by:

$$A_1, A_2, \dots, A_n \vDash C$$

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

Let me be very clear about something. ‘ $\vDash$ ’ is not a symbol of TFL. Rather, it is a symbol of our metalanguage, augmented English (recall the difference between object language and metalanguage from §7). So the metalanguage sentence:

- $P, P \rightarrow Q \models Q$

is *just* an abbreviation for this metalanguage sentence:

- The TFL sentences ' $P$ ' and ' $P \rightarrow Q$ ' tautologically entail ' $Q$ '

Note that there is no limit on the number of TFL sentences that can be mentioned before the symbol ' $\models$ '. Indeed, we can even consider the limiting case:

$$\models C$$

This says that there is no valuation which makes all the sentences mentioned on the left side of ' $\models$ ' true whilst making  $C$  false. Since *no* sentences are mentioned on the left side of ' $\models$ ' in this case, this just means that there is no valuation which makes  $C$  false. Otherwise put, it says that every valuation makes  $C$  true. Otherwise put, it says that  $C$  is a tautology. Equally, to say that  $A$  is a contradiction, we can write:

$$A \models$$

For this says that no valuation makes  $A$  true.

Sometimes, we will want to deny that there is a tautological entailment, and say something of this shape:

$$\text{it is not the case that } A_1, \dots, A_n \models C$$

In that case, we can just slash the turnstile through, and write:

$$A_1, A_2, \dots, A_n \not\models C$$

This means that *some* valuation makes all of  $A_1, \dots, A_n$  true whilst making  $C$  false. (But note that it does *not* immediately follow that  $A_1, \dots, A_n \models \neg C$ , for that would mean that *every* valuation makes all of  $A_1, \dots, A_n$  true whilst making  $C$  false.)

### 11.7 ' $\models$ ' versus ' $\rightarrow$ '

I now want to compare and contrast ' $\models$ ' and ' $\rightarrow$ '.

Observe:  $A \models C$  iff no valuation of the atomic sentences makes  $A$  true and  $C$  false.

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

Combining these two observations, we see that  $A \rightarrow C$  is a tautology iff  $A \models C$ . But there is a really, really important difference between ' $\models$ ' and ' $\rightarrow$ ':

' $\rightarrow$ ' is a sentential connective of TFL.  
' $\models$ ' is a symbol of augmented English.

Indeed, when ' $\rightarrow$ ' is flanked with two TFL sentences, the result is a longer TFL sentence. By contrast, when we use ' $\models$ ', we form a metalinguistic sentence that *mentions* the surrounding TFL sentences.

## 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:

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 assess the following:

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  $A$  and  $B$  are tautologically equivalent. What can you say about  $A \leftrightarrow B$ ?
2. Suppose that  $(A \wedge B) \rightarrow C$  is neither a tautology nor a contradiction. What can you say about this:  $A, B \models C$ ?
3. Suppose that  $A, B$  and  $C$  are jointly tautologically inconsistent. What can you say about this:  $(A \wedge B \wedge C)$ ?
4. Suppose that  $A$  is a contradiction. What can you say about this:  $A, B \therefore C$ ?
5. Suppose that  $C$  is a tautology. What can you say about this:  $A, B \models C$ ?
6. Suppose that  $A$  and  $B$  are tautologically equivalent. What can you say about  $(A \vee B)$ ?
7. Suppose that  $A$  and  $B$  are *not* tautologically equivalent. What can you say about this:  $(A \vee B)$ ?

E. Consider the following principle:

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

Is this principle correct? Explain your answer.

# Truth table shortcuts

With practice, you will quickly become adept at filling out truth tables. In this section, I want to provide (and justify) some shortcuts which will 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

In §11, we saw how to use truth tables to test for validity. In that test, we look for *bad* lines: lines where the premises are all true and the conclusion is false. Now:

- If the conclusion is true on a line, then that line is not bad. (And we don’t need to evaluate anything *else* on that line to confirm this.)
- If any premise is false on a line, then that line is not bad. (And we don’t need to evaluate anything *else* on that line to confirm this.)

With this in mind, we can speed up our tests for validity quite considerably.

Let’s consider how we might test the following:

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

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, after our first stage, we are left with something like this:

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

where the blanks indicate that we won’t bother with any more investigation (since the line is not bad), and the question-marks indicate that we need to keep digging.

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

$J$	$L$	$\neg L \rightarrow (J \vee L)$	$\neg L$	$J$
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 is certainly not bad, because some premise 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 which makes every premise true makes the conclusion true.

It's probably worth illustrating the tactic again. Consider this argument:

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

Again, we start by evaluating the conclusion. Since this is a disjunction, it is true whenever either disjunct is true, so we can speed things along a bit.

A	B	C	$A \vee B$	$\neg(B \wedge C)$	$(A \vee \neg C)$
T	T	T			<b>T</b>
T	T	F			<b>T</b>
T	F	T			<b>T</b>
T	F	F			<b>T</b>
F	T	T	?	?	<b>FF</b>
F	T	F			<b>TT</b>
F	F	T	?	?	<b>FF</b>
F	F	F			<b>TT</b>

We can now ignore all but the two lines where the sentence after the turnstile is false. Evaluating the two sentences on the left of the turnstile, we get:

A	B	C	$A \vee B$	$\neg(B \wedge C)$	$(A \vee \neg C)$
T	T	T			<b>T</b>
T	T	F			<b>T</b>
T	F	T			<b>T</b>
T	F	F			<b>T</b>
F	T	T	<b>T</b>	<b>F T</b>	<b>FF</b>
F	T	F			<b>TT</b>
F	F	T	<b>F</b>		<b>FF</b>
F	F	F			<b>TT</b>

So the entailment holds! And our shortcuts saved us a *lot* of work.

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

## Practice exercises

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

- $\neg B \wedge B$
- $\neg D \vee D$
- $(A \wedge B) \vee (B \wedge A)$
- $\neg[A \rightarrow (B \rightarrow A)]$
- $A \leftrightarrow [A \rightarrow (B \wedge \neg B)]$
- $\neg(A \wedge B) \leftrightarrow A$
- $A \rightarrow (B \vee C)$
- $(A \wedge \neg A) \rightarrow (B \vee C)$
- $(B \wedge D) \leftrightarrow [A \leftrightarrow (A \vee C)]$

# 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:

$S$	$T$	$U$	$W$	$(U \wedge T) \rightarrow (S \wedge W)$
				<b>F</b>

I have only left space for one line, rather than 16, since we are only looking for one line, on which the sentence is false (hence, also, the ‘F’).

The main logical operator 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:

$S$	$T$	$U$	$W$	$(U \wedge T) \rightarrow (S \wedge W)$
				<b>T F F</b>

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

$S$	$T$	$U$	$W$	$(U \wedge T) \rightarrow (S \wedge W)$
<b>T</b>	<b>T</b>			<b>T T T F F</b>

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:

$S$	$T$	$U$	$W$	$(U \wedge T) \rightarrow (S \wedge W)$
<b>F</b>	<b>T</b>	<b>T</b>	<b>F</b>	<b>T T T F F F F</b>

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.

$$\begin{array}{cccc|c} S & T & U & W & (U \wedge T) \rightarrow (S \wedge W) \\ \hline & & & & \mathbf{T} \end{array}$$

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. Making an arbitrary choice, let's make 'U' false; we can then assign any truth value we like to the other atomic sentences.

$$\begin{array}{cccc|c} S & T & U & W & (U \wedge T) \rightarrow (S \wedge W) \\ \hline \mathbf{F} & \mathbf{T} & \mathbf{F} & \mathbf{F} & \mathbf{F} \mathbf{F} \mathbf{T} \mathbf{T} \mathbf{F} \mathbf{F} \mathbf{F} \end{array}$$

**Tautological equivalence.** To show that two sentences are tautologically 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* tautologically 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/entailment.** 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 tautologically 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 tautologically consistent, or jointly tautologically 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.  $A \leftrightarrow B, B \leftrightarrow C \therefore A \leftrightarrow C$

4

# First-order logic

# Building blocks of FOL

14

## 14.1 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 TFL, 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:

$$L, A \therefore F$$

But the truth-table test will now indicate that this is *invalid*. What has gone wrong?

The problem is not that we have made a mistake while symbolising the argument. This is the best symbolisation we can give *in TFL*. The problem lies with TFL 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 TFL are atomic sentences, and TFL 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 *first-order logic*, or *FOL*.

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

First, we have *names*. In FOL, we indicate these with lowercase italic letters. For instance, we might let ‘*b*’ stand for Bertie, or let ‘*i*’ stand for Willard.

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 FOL, we indicate predicates with uppercase italic letters. For instance, we might let the FOL predicate ‘*D*’ symbolise the English predicate ‘\_\_\_\_ is a dog’. Then the expression ‘*Db*’ will be a sentence in FOL, which symbolises the English sentence ‘Bertie is a dog’. Equally, we might let the FOL 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 FOL sentence ‘ $\exists xDx$ ’, which we might read aloud as ‘there is at least one thing,  $x$ , such that  $x$  is a dog’.

That is the general idea. But FOL is significantly more subtle than TFL. So we will come at it slowly.

## 14.2 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 FOL, 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 FOL:

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

These should be thought of along the lines of proper names in English. But with one difference. ‘Tim Button’ is a proper name, but there are several people with this name. (Equally, there are at least two people with the name ‘P.D. Magnus’.) We live with this kind of ambiguity in English, allowing context to individuate the fact that ‘Tim Button’ refers to the lecturer of this course, and not some other guy. In FOL, 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 TFL, we can provide symbolisation keys. These indicate, temporarily, what a name shall pick out. So we might offer:

$e$ : Elsa  
 $g$ : Gregor  
 $m$ : Marybeth

## 14.3 Predicates

The simplest predicates are 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 \_\_\_\_

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

$A$ : \_\_\_\_<sub>1</sub> is angry  
 $H$ : \_\_\_\_<sub>1</sub> is happy

(Why the numerical *subscripts* on the gaps? I shall return to this in §16.)

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 TFL, and symbolise the entire sentence by ‘ $Ag \wedge Am$ ’. This illustrates an important point: FOL has all of the truth-functional connectives of TFL.

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

## 14.4 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. In order to do this, we introduce the ‘ $\forall$ ’ symbol. This is called the UNIVERSAL QUANTIFIER.

A quantifier must always be followed by a variable. In FOL, variables are italic lowercase letters ‘ $s$ ’ through ‘ $z$ ’, with or without subscripts. So 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 and put them in as ‘ $x$ ’. The subsequent ‘ $Hx$ ’ indicates, of that thing you picked out, that it is happy.

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 another new symbol: 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  $x$ ,  $x$  is angry’,  $\exists xAx$  is read naturally as ‘there is something,  $x$ , such that  $x$  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:

6. No one is angry.
7. There is someone who is not happy.
8. Not everyone is happy.

Sentence 6 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 6 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  $\neg\exists xA$ . (Notice that I have here returned to the practice of using ‘ $A$ ’ as a metavariable, from §7.) 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 7 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 have been a perfectly adequate symbolisation of sentence 8.

## 14.5 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 the set of 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 FOL, the domain must always include at least one thing. Moreover, in English we can legitimately infer ‘something is angry’ from ‘Gregor is angry’. In FOL, 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. Every 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.

# Sentences with one quantifier

15

We now have all of the pieces of FOL. Symbolising more complicated sentences is just a matter of knowing how to combine predicates, names, quantifiers, and connectives. There is a knack to this, and there is no substitute for practice.

## 15.1 Dealing with syncategorematic 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:

$W$ : \_\_\_\_<sub>1</sub> is white  
 $C$ : \_\_\_\_<sub>1</sub> is a car  
 $h$ : Herbie

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

2. Damon Stoudamire is a short basketball player.
3. Damon Stoudamire is a man.
4. Damon Stoudamire is a short man.

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

$S$ : \_\_\_\_<sub>1</sub> is short  
 $B$ : \_\_\_\_<sub>1</sub> is a basketball player  
 $M$ : \_\_\_\_<sub>1</sub> is a man  
 $d$ : Damon Stoudamire

Then we would symbolise sentence 2 with ‘ $Sd \wedge Bd$ ’, sentence 3 with ‘ $Md$ ’ and sentence 4 with ‘ $Sd \wedge Md$ ’. But that would be a terrible mistake! For this now suggests that sentences 2 and 3 together *entail* sentence 4. But they do not. Standing at 5’10”, Damon Stoudamire is one of the shortest professional basketball players of all time, but he is nevertheless an averagely-tall man. The point is that sentence 2 says that Damon is short *qua* basketball player, even though he is of average height *qua* man. So you will need to symbolise ‘\_\_\_\_ is a short basketball player’ and ‘\_\_\_\_ is a short man’ using completely different predicates.

Similar examples abound. All footballers are people, but some good footballers are bad people. I might be an incompetent player of the cor anglais, but a competent individual. And so it goes. The moral is: when you see two adjectives in a row, you need to ask yourself carefully whether they can be treated as a conjunction or not.

## 15.2 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

$P$ : \_\_\_<sub>1</sub> is in my pocket

$T$ : \_\_\_<sub>1</sub> is on the table

$Q$ : \_\_\_<sub>1</sub> is a quarter

$D$ : \_\_\_<sub>1</sub> 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 of which ‘ $(Tx \rightarrow Dx)$ ’ is true. Recall that, in TFL,  $A \rightarrow B$  is tautologically equivalent to  $\neg A \vee B$ . This equivalence will also hold in FOL. So ‘ $\exists x(Tx \rightarrow Dx)$ ’ is true if there is some object in the domain, such that ‘ $(\neg Tx \vee Dx)$ ’ is true of that object. That is, ‘ $\exists x(Tx \rightarrow Dx)$ ’ is true if some coin is *either* not on the table *or* is a dime. Of course there is a coin that is not on the table: there are coins lots of other places. So it is *very easy* for ‘ $\exists x(Tx \rightarrow Dx)$ ’ to be true. A conditional will usually be the natural connective to use with a universal quantifier, but a conditional within the scope of an existential quantifier 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  $\neg\forall xA$  and  $\exists x\neg A$ , mentioned in §14, along with the equivalence between  $\neg(A \rightarrow B)$  and  $A \wedge \neg B$ .)

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.

### 15.3 Empty predicates

In §14, 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

$M$ : \_\_\_\_<sub>1</sub> is a monkey.

$S$ : \_\_\_\_<sub>1</sub> 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 that every monkey knows sign language unless 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: the truth of that sentence requires a (signing) monkey.

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

$R$ : \_\_\_\_<sub>1</sub> 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 every refrigerator *in the domain* is also a monkey. Since the domain comprises only *animals*, there are no refrigerators in the domain. Again, then, the sentence is *vacuously* true. (Of course, 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.)

When  $F$  is an empty predicate, any sentence  $\forall x(Fx \rightarrow \dots)$  is vacuously true.

## 15.4 Picking a domain

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

11. Every rose has a thorn.

We might offer this symbolisation key:

$R$ : \_\_\_\_<sub>1</sub> is a rose  
 $T$ : \_\_\_\_<sub>1</sub> has a thorn

It is tempting to say that sentence 11 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 11 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 11 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 FOL. 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  
 $C$ : \_\_\_\_<sub>1</sub> is a cowboy  
 $S$ : \_\_\_\_<sub>1</sub> sings a sad, sad song  
 $R$ : \_\_\_\_<sub>1</sub> is a rose  
 $T$ : \_\_\_\_<sub>1</sub> 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’.

## 15.5 The utility of paraphrase

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

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

domain: people

$B$ : \_\_\_\_<sub>1</sub> is a bassist

$R$ : \_\_\_\_<sub>1</sub> 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 person 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, removing pronouns.

Sentence 13 can be paraphrased as, ‘If Kim Deal is a bassist, then *Kim Deal* is a rockstar’. This can obviously be symbolised as ‘ $Bk \rightarrow Rk$ ’.

Sentence 14 must be paraphrased differently: ‘If a person is a bassist, then *that person* is a rock star’. This sentence is not about any particular person, so we need a variable. As a halfway house, we can paraphrase this as, ‘For any person  $x$ , 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 ‘Everyone 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 she is 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 logical operator: ‘ $\exists xBx \rightarrow Rk$ ’.

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’.

## 15.6 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 need 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 TFL, 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 FOL.

B. Using the following symbolisation key:

domain: people

*K*: \_\_\_\_<sub>1</sub> knows the combination to the safe

*S*: \_\_\_\_<sub>1</sub> is a spy

*V*: \_\_\_\_<sub>1</sub> is a vegetarian

*h*: Hofthor

*i*: Ingmar

symbolise the following sentences in FOL:

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

*A*: \_\_\_\_<sub>1</sub> is an alligator.

*M*: \_\_\_\_<sub>1</sub> is a monkey.

*R*: \_\_\_\_<sub>1</sub> is a reptile.

*Z*: \_\_\_\_<sub>1</sub> lives at the zoo.

*a*: Amos

*b*: Bouncer

*c*: Cleo

symbolise each of the following sentences in FOL:

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 FOL.

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.

# Multiple generality

16

So far, we have only considered sentences that require one-place predicates and one quantifier. The full power of FOL 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.

## 16.1 Many-placed 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.

## 16.2 Mind the gap(s)!

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. These are three perfectly good sentences, obtained by filling in the gaps in ‘\_\_\_\_ loves \_\_\_\_’ in different ways; but they all have distinctively different meanings:

1. Karl loves Imre.
2. Imre loves Karl.
3. Karl loves Karl.

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 assign them numbers. And my method is explained by example. Suppose I want to symbolise the preceding sentences. Then I might start with the following representation key:

domain: people  
*i*: Imre  
*k*: Karl  
*L*: \_\_\_\_<sub>1</sub> loves \_\_\_\_<sub>2</sub>

Sentence 1 will be symbolised by ‘*Lki*’, sentence 2 will be symbolised by ‘*Lik*’, and sentence 3 will be symbolised by ‘*Lkk*’. Here are a few more sentences that I can symbolise with the same key:

4. Imre loves himself.
5. Karl loves Imre, but not vice versa.
6. Karl is loved by Imre.

Sentence 4 can be paraphrased as ‘Imre loves Imre’, and so symbolised by ‘*Lii*’. Sentence 5 is a conjunction. We can paraphrase it as ‘Karl loves Imre, and Imre does not love Karl’, and so symbolise it as ‘*Lki* ∧ ¬*Lik*’. Sentence 6 can be paraphrased by ‘Imre loves Karl’, and so symbolised as ‘*Lik*’. In this last case, of course, we have lots the difference in *tone* between the active and passive voice; but we have at least preserved the truth conditions.

But the relationship between ‘Imre loves Karl’ and ‘Karl is loved by Imre’ highlights something important. To see what, suppose we add another entry to our symbolisation key:

*M*: \_\_\_\_<sub>2</sub> loves \_\_\_\_<sub>1</sub>

The entry for ‘*M*’ uses exactly the same English word—‘loves’—as the entry for ‘*L*’. *But the gaps have been swapped around!* (Just look closely at the subscripts.) And this *matters*.

To explain: when we see a sentence like ‘*Lki*’, we are being told to take the *first* name (i.e. ‘*k*’) and associate its value (i.e. Karl) with the gap labelled ‘1’, then take the *second* name (i.e. ‘*i*’) and associate its value (i.e. Imre) with the gap labelled ‘2’, and so come up with: *Karl loves Imre*. The sentence ‘*Mik*’ also tells us to take the *first* name (i.e. ‘*i*’) and plug its value into the gap labelled ‘1’, and take the *second* name (i.e. ‘*k*’) and plug its value into the gap labelled ‘2’, and so come up with: *Imre loves Karl*.

So, ‘*Lik*’ and ‘*Mki*’ both symbolise ‘Imre loves Karl’, whereas ‘*Lki*’ and ‘*Mik*’ both symbolise ‘Karl loves Imre’. Since love can be unrequited, these are different claims.

One last example might be helpful. Suppose we add this to our symbolisation key:

*P*: \_\_\_\_<sub>1</sub> prefers \_\_\_\_<sub>1</sub> to \_\_\_\_<sub>2</sub>

Now the sentence ‘*Pik*’ symbolises ‘Imre prefers Imre to Karl’, and ‘*Pki*’ symbolises ‘Karl prefers Karl to Imre’. And note that we could have achieved the same effect, if we had instead specified:

*P*: \_\_\_\_<sub>1</sub> prefers themselves to \_\_\_\_<sub>2</sub>

In any case, the overall moral of this is simple. *When dealing with predicates with more than one place, pay careful attention to the order of the gaps!*

### 16.3 The order of quantifiers

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

7. For every person, there is some person whom the former loves
8. There is some particular person whom every person loves

Sentence 7 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 7 will be true.

Sentence 8 can be symbolised by ‘ $\exists y\forall xLxy$ ’. Sentence 8 is *not* true in the situation just described. Again, suppose that our domain of discourse is restricted to Imre, Juan and Karl. Then all of Juan, Imre and Karl must converge on (at least) one object of love.

The point of the example is to illustrate that the order 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. (∀∃)

So: There is some particular truth that no person can know. (∃∀)

This argument form is obviously invalid. It’s just as bad as:<sup>1</sup>

Every dog has its day. (∀∃)

So: There is a day for all the dogs. (∃∀)

The moral is simple. *Take care with the order of your quantifiers!*

### 16.4 Stepping-stones to symbolisation

As we are starting to see, symbolisation in FOL can become a bit tricky. So, when symbolising a complex sentence, I recommend laying down several stepping-stones. As usual, the idea is best illustrated by example. Consider this symbolisation key:

domain: people and dogs

*D*: \_\_\_\_<sub>1</sub> is a dog

*F*: \_\_\_\_<sub>1</sub> is a friend of \_\_\_\_<sub>2</sub>

*O*: \_\_\_\_<sub>1</sub> owns \_\_\_\_<sub>2</sub>

*g*: Geraldo

<sup>1</sup>Thanks to Rob Trueman for the example.

And now let's try to symbolise these sentences:

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

Sentence 9 can be paraphrased as, 'There is a dog that Geraldo owns'. This can be symbolised by ' $\exists x(Dx \wedge Ogx)$ '.

Sentence 10 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 9, except that it is not specifically about Geraldo. So we can symbolise sentence 10 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 FOL sentence *nor* an English sentence: it uses bits of FOL (' $\exists$ ', 'y') and bits of English ('dog owner'). It is really is *just a stepping-stone* on the way to symbolising the entire English sentence with a FOL 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 11 can be paraphrased as, 'Everyone who is a friend of Geraldo is a dog owner'. Using our stepping-stone tactic, we might write

$$\forall x [Fyg \rightarrow x \text{ is a dog owner}]$$

Now the fragment that we have left to deal with, 'x is a dog owner', is structurally just like sentence 9. But it would be a mistake for us simply to write

$$\forall x [Fyg \rightarrow \exists x (Dx \wedge Oxx)]$$

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. And now confusion reigns: which ' $x$ ' are we talking about? Suddenly the sentence would be ambiguous (if it is even meaningful at all), and logicians hate ambiguity. The broad moral is that a single variable cannot serve two masters simultaneously.

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

$$\forall x [Fyg \rightarrow \exists z (Dz \wedge Oxz)]$$

and this adequately symbolises sentence 11.

Sentence 12 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 clash here, because it is clear which quantifier the letter falls under. We might graphically represent the scope of the quantifiers thus:

$$\begin{array}{c} \text{scope of '}\forall x\text{' } \\ \overbrace{\hspace{15em}} \\ \text{scope of '}\exists y\text{' } \\ \overbrace{\hspace{10em}} \\ \text{scope of 1st '}\exists z\text{' } \quad \text{scope of 2nd '}\exists z\text{' } \\ \overbrace{\hspace{5em}} \quad \overbrace{\hspace{5em}} \\ \forall x[\exists z(Dz \wedge Oxz) \rightarrow \exists y(\exists z(Dz \wedge Oyz) \wedge Fyx)] \end{array}$$

This shows that no variable is being forced to serve two masters simultaneously.

Sentence 13 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

A: \_\_\_\_\_<sub>1</sub> is an alligator

M: \_\_\_\_\_<sub>1</sub> is a monkey

R: \_\_\_\_\_<sub>1</sub> is a reptile

Z: \_\_\_\_\_<sub>1</sub> lives at the zoo

L: \_\_\_\_\_<sub>1</sub> loves \_\_\_\_\_<sub>2</sub>

a: Amos

b: Bouncer

c: Cleo

symbolise each of the following sentences in FOL:

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

*D*: \_\_\_\_<sub>1</sub> is a dog

*S*: \_\_\_\_<sub>1</sub> likes samurai movies

*L*: \_\_\_\_<sub>1</sub> is larger than \_\_\_\_<sub>2</sub>

*b*: Bertie

*e*: Emerson

*f*: Fergis

symbolise the following sentences in FOL:

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

*R*: \_\_\_\_<sub>1</sub> has run out.

*T*: \_\_\_\_<sub>1</sub> is on the table.

*F*: \_\_\_\_<sub>1</sub> is food.

*P*: \_\_\_\_<sub>1</sub> is a person.

*L*: \_\_\_\_<sub>1</sub> likes \_\_\_\_<sub>2</sub>.

*e*: Eli

*f*: Francesca

*g*: the guacamole

symbolise the following English sentences in FOL:

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

*D*: \_\_\_\_<sub>1</sub> dances ballet.

*F*: \_\_\_\_<sub>1</sub> is female.

*M*: \_\_\_\_<sub>1</sub> is male.

*C*: \_\_\_\_<sub>1</sub> is a child of \_\_\_\_<sub>2</sub>.

*S*: \_\_\_\_<sub>1</sub> is a sibling of \_\_\_\_<sub>2</sub>.

*e*: Elmer

*j*: Jane

*p*: Patrick

symbolise the following arguments in FOL:

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.

# Identity

17

Consider this sentence:

1. Pavel owes money to everyone

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

$O$ : \_\_\_\_<sub>1</sub> owes money to \_\_\_\_<sub>2</sub>  
 $p$ : Pavel

we can symbolise sentence 1 by ‘ $\forall xOp_x$ ’. 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. And maybe we did not want to say that. Maybe 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*

But we do not have any way for dealing with the italicised words yet. The solution is to add another symbol to FOL.

## 17.1 Adding identity

The symbol ‘=’ will be a two-place predicate. Since it will have a special meaning, we shall write it a bit differently: we put it *between* two terms, rather than out front. (This should also be familiar; consider a mathematical equation like  $\frac{1}{2} = 0.5$ .) And the special meaning for ‘=’ is given by the fact that we *always* adopt the following symbolisation key:

$=$ : \_\_\_\_<sub>1</sub> is identical to \_\_\_\_<sub>2</sub>

This 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.

To put this to use, suppose we want to symbolise this sentence:

5. Pavel is Mister Checkov.

Let us add to our symbolisation key:

$c$ : Mister Checkov

Now sentence 5 can be symbolised as ' $p = c$ '. This tells us that the names ' $p$ ' and ' $c$ ' both name the very same thing.

We can also now deal with sentences 2–4. All of these sentences can be paraphrased as 'Everyone who is not Pavel is owed money by Pavel'. 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, ' $x = p$ '.

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:

6. No one besides Pavel owes money to Hikaru.
7. Only Pavel owes Hikaru money.

Letting ' $h$ ' name Hikaru, sentence 6 can be paraphrased as, 'No one who is not Pavel owes money to Hikaru'. This can be symbolised by ' $\neg \exists x(\neg x = p \wedge Oxh)$ '. Equally, sentence 6 can be paraphrased as 'for all  $x$ , if  $x$  owes money to Hikaru, then  $x$  is Pavel'. Then it can be symbolised as ' $\forall x(Oxh \rightarrow x = p)$ '.

Sentence 7 can be treated similarly. But there is one subtlety here. Do either sentence 6 or 7 entail that Pavel himself owes money to Hikaru?

## 17.2 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:

8. There is at least one apple
9. There are at least two apples
10. There are at least three apples

We shall use the symbolisation key:

A: \_\_\_\_<sub>1</sub> is an apple

Sentence 8 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 9 without identity. But 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. To make sure that we are dealing with *different* apples, we need to use identity. Sentence 9 needs to say that the apples are distinct, so it can be symbolised by ' $\exists x\exists y(Ax \wedge Ay \wedge \neg x = y)$ '.

Sentence 10 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: ' $\exists x\exists y\exists z(Ax \wedge Ay \wedge Az \wedge \neg x = y \wedge \neg y = z \wedge \neg x = z)$ '.

### 17.3 There are at most...

Now consider these sentences:

11. There is at most one apple
12. There are at most two apples

Sentence 11 can be paraphrased as, ‘It is not the case that there are at least *two* apples’. This is just the negation of sentence 9:

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

But sentence 11 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.

Similarly, sentence 12 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:

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

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:

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

### 17.4 There are exactly...

We can now consider precise statements of numerical quantity, like:

13. There is exactly one apple.
14. There are exactly two apples.
15. There are exactly three apples.

Sentence 13 can be paraphrased as, ‘There is *at least* one apple and there is *at most* one apple’. This is just the conjunction of sentence 8 and sentence 11. 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 13 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 14 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:

16. There are exactly two things
17. 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:

$$\begin{aligned} \exists x \exists y \neg x = y \wedge \neg \exists x \exists y \exists z (\neg x = y \wedge \neg y = z \wedge \neg x = z) \\ \exists x \exists y [\neg x = y \wedge \forall z (x = z \vee y = z)] \end{aligned}$$

## 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’.

# Definite descriptions

18

Consider sentences like:

1. Nick is the traitor.
2. The traitor went to Cambridge.
3. The traitor is the deputy

These are definite descriptions: they are meant to pick out a *unique* object. They should be contrasted with *indefinite* descriptions, such as ‘Nick is *a* traitor’. They should equally be contrasted with *generics*, such as ‘*The* whale is a mammal’ (when it’s inappropriate to ask *which* whale). The question we face is: how should we deal with definite descriptions in FOL?

## 18.1 Treating definite descriptions as terms

One option would be to introduce new names whenever we come across a definite description. This is probably not a great idea. We know that *the* traitor—whoever it is—is indeed *a* traitor. We want to preserve that information in our symbolisation.

A second option would be to use a *new* definite description operator, such as ‘*t*’. The idea would be to symbolise ‘the F’ as ‘ $\iota xFx$ ’ (think ‘the *x* such that *Fx*’); or to symbolise ‘the G’ as ‘ $\iota xGx$ ’, etc. Expressions of the form  $\iota xAx$  would then behave like names. If we were to follow this path, then we could use the following symbolisation key:

domain: people

*T*: \_\_\_\_<sub>1</sub> is a traitor

*D*: \_\_\_\_<sub>1</sub> is a deputy

*C*: \_\_\_\_<sub>1</sub> went to Cambridge

*n*: Nick

Then, we could symbolise sentence 1 with ‘ $\iota xTx = n$ ’, sentence 2 with ‘ $C\iota xTx$ ’, and sentence 3 with ‘ $\iota xTx = \iota xDx$ ’.

However, it would be nice if we didn’t have to add a new symbol to FOL. And we might be able to make do without one.

## 18.2 Russell’s analysis

Bertrand Russell offered an *analysis* of definite descriptions. Very briefly put, he observed that, when we say ‘the F’ in the context of a definite description, our aim is to pick out the *one and only* thing that is F (in the appropriate context). Thus Russell analysed the notion

of a definite description as follows:<sup>1</sup>

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 analysis: *‘the’ does not appear on the right-side of the equivalence*. Russell is aiming to provide an understanding of definite descriptions, in terms that do not presuppose them.

Now, one might worry that I can say ‘the table is brown’ without implying that there is one and only one table in the universe. But this is not (yet) a fantastic counterexample to Russell’s analysis. The domain of discourse is likely to be restricted by context (e.g. to salient objects in my vicinity).

If we accept Russell’s analysis of definite descriptions, then we can symbolise sentences of the form ‘the F is G’ using our strategy for numerical quantification in FOL. After all, we can deal with the three conjuncts on the right-hand side of Russell’s analysis 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 G. 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 ‘*r*’.

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 analysis, 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:

$$\exists x\exists y([\exists z(Tz \wedge \forall z(Tz \rightarrow x = z))] \wedge [\exists z(Dz \wedge \forall z(Dz \rightarrow y = z))] \wedge x = y)$$

Note that the formula ‘ $x = y$ ’ must fall within the scope of both quantifiers!

### 18.3 Empty definite descriptions

One of the nice features of Russell’s analysis 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 §14 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.

<sup>1</sup>Bertrand Russell, ‘On Denoting’, 1905, *Mind* 14, pp. 479–93; also Russell, *Introduction to Mathematical Philosophy*, 1919, London: Allen and Unwin, ch. 16.

Russell's analysis neatly avoids this problem. Russell tells us to treat definite descriptions using predicates and quantifiers, instead of names. Since predicates can be empty (see §15), this means that no difficulty now arises when the definite description is empty.

Indeed, Russell's analysis helpfully highlights two ways to go wrong in a claim involving a definite description. To adapt an example from Stephen Neale (1990),<sup>2</sup> suppose I, Tim Button, claim:

4. The present king of France is my spouse

Using the following symbolisation key:

*b*: Tim

*K*: \_\_\_\_<sub>1</sub> is a present king of France

*M*: \_\_\_\_<sub>1</sub> is married to \_\_\_\_<sub>2</sub>

(NB: the symbolisation key speaks of *a* present King of France, not *the* present King of France; i.e. it employs indefinite, rather than definite, description.) Sentence 4 would be symbolised by ' $\exists x(\forall y(Ky \leftrightarrow x = y) \wedge Bxb)$ '. Now, this can be false in (at least) two ways, corresponding to these two different sentences:

5. There is no one who is both the present King of France and Tim's spouse.  
6. There is a unique present King of France, but he is not Tim's spouse.

Sentence 5 can be symbolised by ' $\neg\exists x[Kx \wedge \forall y(Ky \rightarrow x = y) \wedge Mbx]$ '. We might call this *outer* negation, since the negation governs the entire sentence. Note that the sentence is true if there is no present King of France.

Sentence 6 can be symbolised by ' $\exists x(Kx \wedge \forall y(Ky \rightarrow x = y) \wedge \neg Mbx)$ '. 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 not Tim's spouse.

## 18.4 The adequacy of Russell's analysis

How good is Russell's analysis of definite descriptions? This question has generated a substantial philosophical literature, 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 analysis, both 'the F is G' is and 'the F is non-G' are false. P.F. Strawson suggested that such sentences should not be regarded as false, exactly, but involve *presupposition failure*, and so need to be treated as *neither true nor false*.<sup>3</sup>

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 sentence is assigned exactly one of these truth values.

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 the present King of France is my spouse?<sup>4</sup>

Keith Donnellan raised a second sort of worry, which (very roughly) can be brought out by thinking about a case of mistaken identity.<sup>5</sup> Two men stand in the corner: a very tall

<sup>2</sup>Neale, *Descriptions*, 1990, Cambridge: MIT Press.

<sup>3</sup>P.F. Strawson, 'On Referring', 1950, *Mind* 59, pp. 320–34.

<sup>4</sup>This is Neale's (1990) line.

<sup>5</sup>Keith Donnellan, 'Reference and Definite Descriptions', 1966, *Philosophical Review* 77, pp. 281–304.

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:

7. The gin-drinker is very tall!

Russell's analysis will have us render Malika's sentence as:

7'. 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 analysis, 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 analysis, she actually picked out a different man (the short one), and consequently said something false of him. But maybe advocates of Russell's analysis 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.<sup>6</sup>

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 analysis of definite descriptions, when it comes to putting things into FOL. It is certainly the best that we can offer, without significantly revising our logic. And it is quite defensible as an analysis.

## Practice exercises

A. Using the following symbolisation key:

domain: people

*K*: \_\_\_\_<sub>1</sub> knows the combination to the safe.

*S*: \_\_\_\_<sub>1</sub> is a spy.

*V*: \_\_\_\_<sub>1</sub> is a vegetarian.

*T*: \_\_\_\_<sub>1</sub> trusts \_\_\_\_<sub>2</sub>.

*h*: Hofthor

*i*: Ingmar

symbolise the following sentences in FOL:

1. Hofthor 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 Hofthor, 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.

<sup>6</sup>For more, see Saul Kripke, 'Speaker Reference and Semantic Reference', 1977.

B. Using the following symbolisation key:

domain: cards in a standard deck

$B$ :  $\_\_\_\_1$  is black.

$C$ :  $\_\_\_\_1$  is a club.

$D$ :  $\_\_\_\_1$  is a deuce.

$J$ :  $\_\_\_\_1$  is a jack.

$M$ :  $\_\_\_\_1$  is a man with an axe.

$O$ :  $\_\_\_\_1$  is one-eyed.

$W$ :  $\_\_\_\_1$  is wild.

symbolise each sentence in FOL:

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

$B$ :  $\_\_\_\_1$  is in Farmer Brown's field.

$H$ :  $\_\_\_\_1$  is a horse.

$P$ :  $\_\_\_\_1$  is a Pegasus.

$W$ :  $\_\_\_\_1$  has wings.

symbolise the following sentences in FOL:

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 ' $\exists x(Tx \wedge \forall y(Ty \rightarrow x = y) \wedge x = n)$ '. Explain why these would be equally good symbolisations:

- $Tn \wedge \forall y(Ty \rightarrow n = y)$
- $\forall y(Ty \leftrightarrow y = n)$

# Sentences of FOL

19

We know how to represent English sentences in FOL. The time has finally come to define the notion of a *sentence* of FOL.

## 19.1 Expressions

There are six kinds of symbols in FOL:

Predicates with subscripts, as needed	$A, B, C, \dots, Z$ $A_1, B_1, Z_1, A_2, A_{25}, f_{375}, \dots$
Constants with subscripts, as needed	$a, b, c, \dots, r$ $a_1, b_{224}, h_7, m_{32}, \dots$
Variables with subscripts, as needed	$s, t, u, v, w, x, y, z$ $x_1, y_1, z_1, x_2, \dots$
Connectives	$\neg, \wedge, \vee, \rightarrow, \leftrightarrow$
Brackets	$(, )$
Quantifiers	$\forall, \exists$

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

## 19.2 Terms and formulas

In §6, we went straight from the statement of the vocabulary of TFL to the definition of a sentence of TFL. In FOL, 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 intuitive idea will take some time to unpack.

We start by defining the notion of a term.

A TERM is any name or any variable.

So, here are some terms:

$$a, b, x, x_1 x_2, y, y_{254}, z$$

We next need to define atomic formulas.

1. If  $\mathbf{R}$  is an  $n$ -place predicate and  $\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_n$  are terms, then  $\mathbf{R}\mathbf{t}_1\mathbf{t}_2\dots\mathbf{t}_n$  is an atomic formula.
2. If  $\mathbf{t}_1$  and  $\mathbf{t}_2$  are terms, then  $\mathbf{t}_1 = \mathbf{t}_2$  is an atomic formula.
3. Nothing else is an atomic formula.

This use of bold-italic-fonts follows the conventions laid down in §7. So, ' $\mathbf{R}$ ' is not itself a predicate of FOL. Rather, it is a symbol of our metalanguage (augmented English) which I am using here to talk about any predicate of FOL. Similarly, ' $\mathbf{t}_1$ ' is not a term of FOL, but a symbol of the metalanguage that I am using to talk about any term of FOL.

Now, let ' $F$ ' be a one-place predicate, let ' $G$ ' is a three-place predicate, and let ' $S$ ' is a six-place predicate. Here are some atomic formulas:

$$\begin{aligned} &x = a \\ &a = b \\ &Fx \\ &Fa \\ &>Gxay \\ &Gaaa \\ &Sx_1x_2abyx_1 \\ &Sby_{254}zaaz \end{aligned}$$

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 TFL.

1. Every atomic formula is a formula.
2. If  $\mathbf{A}$  is a formula, then  $\neg\mathbf{A}$  is a formula.
3. If  $\mathbf{A}$  and  $\mathbf{B}$  are formulas, then  $(\mathbf{A} \wedge \mathbf{B})$  is a formula.
4. If  $\mathbf{A}$  and  $\mathbf{B}$  are formulas, then  $(\mathbf{A} \vee \mathbf{B})$  is a formula.
5. If  $\mathbf{A}$  and  $\mathbf{B}$  are formulas, then  $(\mathbf{A} \rightarrow \mathbf{B})$  is a formula.
6. If  $\mathbf{A}$  and  $\mathbf{B}$  are formulas, then  $(\mathbf{A} \leftrightarrow \mathbf{B})$  is a formula.
7. If  $\mathbf{A}$  is a formula,  $\mathbf{x}$  is a variable,  $\mathbf{A}$  contains at least one occurrence of  $\mathbf{x}$ , and  $\mathbf{A}$  contains neither  $\forall\mathbf{x}$  nor  $\exists\mathbf{x}$ , then  $\forall\mathbf{x}\mathbf{A}$  is a formula.
8. If  $\mathbf{A}$  is a formula,  $\mathbf{x}$  is a variable,  $\mathbf{A}$  contains at least one occurrence of  $\mathbf{x}$ , and  $\mathbf{A}$  contains neither  $\forall\mathbf{x}$  nor  $\exists\mathbf{x}$ , then  $\exists\mathbf{x}\mathbf{A}$  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:

$$\begin{aligned} &Fx \\ &Gayz \\ &Syzyayx \\ &(Gayz \rightarrow Syzyayx) \\ &\forall z(Gayz \rightarrow Syzyayx) \end{aligned}$$

$$\begin{aligned}
 &Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx) \\
 &\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx)) \\
 &\forall x\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx))
 \end{aligned}$$

But this is *not* a formula:

$$\forall x\exists xGxxx$$

Certainly ‘ $Gxxx$ ’ is a formula. And ‘ $\exists xGxxx$ ’ is therefore also a formula. But we cannot form a new formula by putting ‘ $\forall x$ ’ at the front. This violates the constraints on clause 7 of our recursive definition: ‘ $\exists xGxxx$ ’ contains at least one occurrence of ‘ $x$ ’, but it already contains ‘ $\exists x$ ’.

These constraints have the effect of ensuring that variables only serve one master at any one time (see §16). And in fact, 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 TFL, though we note that a logical operator can be either a connective or a quantifier:

The MAIN LOGICAL OPERATOR in a formula is the operator that was introduced most recently, when that formula was constructed using the recursion rules.

The SCOPE of a logical operator in a formula is the subformula for which that operator is the main logical operator.

So we can graphically illustrate the scope of the quantifiers in the preceding example thus:

$$\begin{array}{c}
 \text{scope of '}\forall x\text{' } \\
 \overbrace{\hspace{10em}} \\
 \text{scope of '}\exists y\text{' } \\
 \overbrace{\hspace{8em}} \\
 \text{scope of '}\forall z\text{' } \\
 \overbrace{\hspace{6em}} \\
 \forall x\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx))
 \end{array}$$

### 19.3 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

$L$ : \_\_\_\_<sub>1</sub> loves \_\_\_\_<sub>2</sub>

$b$ : 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.

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.

An occurrence of a variable  $x$  is BOUND iff it falls within the scope of either  $\forall x$  or  $\exists x$ . A variable which is not bound is FREE.

For example, consider the formula

$$(\forall x(Ex \vee Dy) \rightarrow \exists z(Ex \rightarrow Lzx))$$

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 ' $(Ex \rightarrow Lzx)$ ', so ' $z$ ' is bound.

Finally we can say the following.

A SENTENCE of FOL is any formula of FOL that contains no free variables.

## 19.4 Bracketing conventions

We will adopt the same notational conventions governing brackets that we did for TFL (see §6 and §10.3.)

First, we may omit the outermost brackets of a formula.

Second, we may use square brackets, '[' and ']', in place of brackets to increase the readability of formulas.

Third, we may omit brackets between each pair of conjuncts when writing long series of conjunctions.

Fourth, we may omit brackets between each pair of disjuncts when writing long series of disjunctions.

## 19.5 Superscripts on predicates

Above, I said that an  $n$ -place predicate followed by  $n$  terms is an atomic formula. But there is a small issue with this definition: the symbols we use for predicates do not, themselves, indicate how many places the predicate has. Indeed, in some places in this book, I have used the letter ' $G$ ' as a one-place predicate; in other places I have used it as a three-place predicate. So, unless I state explicitly whether I want to use ' $G$ ' as a one-place predicate or as a three place predicate, it is *indeterminate* whether ' $Ga$ ' is an atomic formula or not.

There is an easy way to avoid this, which many books adopt. Instead of saying that our predicates are just capital letters (with numerical subscripts as necessary), we could say that they are capital letters *with numerical superscripts* (and with numerical subscripts as necessary). The purpose of the superscript would be to say explicitly how many places the predicate has. On this approach, ' $G^1$ ' would be a one-place predicate, and ' $G^3$ ' would be an (entirely different) three places predicate. They would need to have different entries in any symbolisation key. And ' $G^1a$ ' would be an atomic formula, whereas ' $G^3a$ ' would not; likewise ' $G^3abc$ ' would be an atomic formula, and ' $G^1abc$ ' would not.

So, as I say, we *could* add superscripts to all our predicates. This would have the advantage of making certain things completely explicit. However, it would have the disadvantage of making our formulas much harder to read; the superscripts would distract the eye. So, we will not bother to make this change. Our predicates will remain *without* superscripts. (And, in practice, any book which includes superscripts almost immediately stops including them!)

However, this leaves open a possibility of ambiguity. So, when any ambiguity could arise—in practice, very rarely—you should say, explicitly, how many places your predicate(s) have.

### 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.  $\forall x\exists y[Rxy \rightarrow (\exists z \wedge Kx)] \vee Ryx$
5.  $\forall x_1(Mx_2 \leftrightarrow Lx_2x_1) \wedge \exists x_2Lx_3x_2$

5

## Interpretations

Recall that TFL is a truth-functional language. Its connectives are all truth-functional, and *all* that we can do with TFL is key sentences to particular truth values. We can do this *directly*. For example, we might stipulate that the TFL 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 is *just* a means of specifying ‘ $P$ ’s truth value; the symbolisation key statement amounts to something like the following stipulation:

- the TFL sentence ‘ $P$  is true iff Big Ben is in London

And I emphasised in §9 that TFL cannot handle differences in meaning that go beyond mere differences in truth value.

## 20.1 Symbolising versus translating

FOL has some similar limitations. It gets beyond mere truth values, since it enables us to split up sentences into terms, predicates and quantifiers. This enables us to consider what is *true of* some particular object, or of some or all objects. *But that’s it*.

To unpack this a bit, consider this symbolisation key:

$C$ : \_\_\_\_<sub>1</sub> lectures logic in Cambridge in Michaelmas 2018

This stipulation does not carry the *meaning* of the English predicate across into our FOL predicate. We are simply stipulating something like this:

- ‘ $C$  is to be true of exactly those things which lecture logic in Cambridge in Michaelmas 2018 (whatever those things might be).

This is an indirect way of stipulating which things a predicate is true of. Alternatively, we can stipulate predicate extensions directly. We can stipulate that ‘ $C$  is to be true of Tim Button, and Tim Button alone. As it happens, this direct stipulation would have the same effect as the indirect stipulation, since Tim, and Tim alone, lectures logic in Cambridge in Michaelmas 2018. But the English predicates ‘\_\_\_\_ is Tim Button’ and ‘\_\_\_\_ lectures logic in Cambridge in Michaelmas 2018’ have very different meanings!

The point is that FOL has no resources for dealing with nuances of meaning. When we interpret FOL, all we are considering is what the predicates are true of. This is normally summed up by saying that FOL is an EXTENSIONAL LANGUAGE.

This is why I speak of *symbolising* English sentences in FOL. It is doubtful that we are *translating* English into FOL, for translation should preserve meaning.

## 20.2 A word on extensions

We can stipulate directly what predicates are to be true of. And our stipulations can be as arbitrary as we like. For example, we could stipulate that ‘ $H$ ’ should be true of, and only of, the following objects:

Theresa May  
the number  $\pi$   
every top-F key on every piano ever made

Armed with this interpretation of ‘ $H$ ’, suppose I now add to my symbolisation key:

$c$ : Jeremy Corbyn  
 $e$ : Theresa May  
 $p$ : the number  $\pi$

Then ‘ $He$ ’ and ‘ $Hp$ ’ will both be true, on this interpretation, but ‘ $Hc$ ’ will be false, since Jeremy Corbyn was not among the stipulated objects.

This process of explicit stipulation is sometimes described as stipulating the *extension* of a predicate. Note that, in the stipulation I just gave, the objects I listed have nothing particularly in common. This doesn’t matter. Logic doesn’t care about what we humans (at a particular moment at spacetime) think ‘naturally go together’; to logic, all objects are on an equal footing.

## 20.3 Many-place predicates

All of this is quite easy to understand when it comes to one-place predicates. But it gets messier when we consider two-place predicates. Consider a symbolisation key like:

$L$ : \_\_\_\_<sub>1</sub> loves \_\_\_\_<sub>2</sub>

Given what I said above, this symbolisation key should be read as saying something like:<sup>1</sup>

- ‘ $L$ ’ is to be true of  $x$  and  $y$  (in that order) iff  $x$  loves  $y$ .

But it is important to insist upon the clause ‘in that order’, since—famously—love is not always reciprocated.

That is an indirect stipulation. What about a direct stipulation? This is also tricky. If we *simply* list objects that fall under ‘ $L$ ’, 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 ‘ $B$ ’ is to be true of, and only of, the following pairs of objects:

⟨Lenin, Marx⟩  
⟨Heidegger, Sartre⟩  
⟨Sartre, Heidegger⟩

<sup>1</sup>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 FOL, and they are being *mentioned*.

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  $\langle \text{Lenin, Marx} \rangle$  is in my explicit list. But '*Bml*' will be false, since  $\langle \text{Marx, Lenin} \rangle$  is not in my list. However, both '*Bhs*' and '*Bsh*' will be true, since both  $\langle \text{Heidegger, Sartre} \rangle$  and  $\langle \text{Sartre, Heidegger} \rangle$  are in my list

To make these ideas more precise, we would need to develop some very elementary *set theory*. Set theory has formal apparatus which allows us to deal with extensions, ordered pairs, and so forth. However, set theory is not covered in this book. So I shall leave these ideas at an imprecise level. I hope that the general idea is clear enough.

## 20.4 Semantics for identity

Identity is a special predicate of FOL. We write it a bit differently than other two-place predicates: ' $x = y$ ' instead of '*Ixy*' (for example). More important, though, its interpretation is fixed, once and for all.

If two names refer to the same object, then swapping one name for another will not change the truth value of any sentence. So, in particular, if '*a*' and '*b*' name the same object, then all of the following will be true:

$$\begin{aligned} Aa &\leftrightarrow Ab \\ Ba &\leftrightarrow Bb \\ Raa &\leftrightarrow Rbb \\ Raa &\leftrightarrow Rab \\ Rca &\leftrightarrow Rcb \\ \forall xRxa &\leftrightarrow \forall xRxb \end{aligned}$$

Some philosophers have believed the reverse of this claim. That is, they have believed that when exactly the same sentences (not containing '=' ) are true of *a* and *b*, then *a* and *b* are the very same object. This is a highly controversial philosophical claim—sometimes called the *identity of indiscernibles*—and our logic will not subscribe to it. So we allow that exactly the same things might be true of two *distinct* objects.

To bring this out, consider the following interpretation:

domain: P.D. Magnus, Tim Button

*a*: P.D. Magnus

*b*: Tim Button

- For every primitive predicate we care to consider, that predicate is true of *nothing*.

Suppose '*A*' is a one-place predicate; then '*Aa*' is false and '*Ab*' is false, so ' $Aa \leftrightarrow Ab$ ' is true. Similarly, if '*R*' is a two-place predicate, then '*Raa*' is false and '*Rab*' is false, so that ' $Raa \leftrightarrow Rab$ ' is true. And so it goes: every atomic sentence not involving '=' is false, so every biconditional linking such sentences is true. For all that, Tim Button and P.D. Magnus are two distinct people, not one and the same!

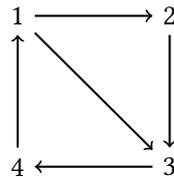
## 20.5 Interpretations

I defined a *valuation* in TFL as any assignment of truth and falsity to atomic sentences. In FOL, 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
- for each predicate that we care to consider (apart from '='), a specification of what things (in what order) the predicate is to be true of. (We don't need to specify an interpretation of '=', since it has a *fixed* interpretation.)

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 in this chapter.

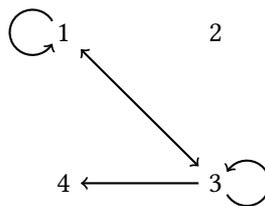
However, it is sometimes also convenient to present an interpretation *diagrammatically*. To illustrate (literally): suppose we want to consider just a single two-place predicate, 'R'. Then we can represent it just by drawing an arrow between two objects, and stipulate that 'R' is to hold of x and y just in case there is an arrow running from x to y in our diagram. As an example, we might offer:



This diagram could be used to describe an interpretation whose domain is the first four positive whole numbers, and which interprets 'R' 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 this diagram:



for an interpretation with the same domain, which takes the extension of 'R' as:

$$\langle 1, 3 \rangle, \langle 3, 1 \rangle, \langle 3, 4 \rangle, \langle 1, 1 \rangle, \langle 3, 3 \rangle$$

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 one-place predicate 'H', say.

I have introduced you to interpretations. 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 I now need to say, precisely, what it is for an arbitrary FOL sentence to be true or false in an interpretation.

We know from §19 that there are three kinds of sentence in FOL:

- atomic sentences
- sentences whose main logical operator is a sentential connective
- sentences whose main logical operator is a quantifier

So I 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

$d$ : Donald Trump

$W$ : \_\_\_\_<sub>1</sub> is wise

$R$ : \_\_\_\_<sub>1</sub> was born before \_\_\_\_<sub>2</sub>

This will be my *go-to example* in what follows.

## 21.1 Atomic sentences

The truth of atomic sentences should be fairly straightforward. The sentence ' $Wa$ ' should be true just in case ' $W$ ' 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, ' $Wd$ ' is (spectacularly) false on our go-to interpretation.

Likewise, on this interpretation, ' $Rad$ ' is true iff the object named by ' $a$ ' was born before the object named by ' $d$ '. Well, Aristotle was born before Donald Trump. So ' $Rad$ ' is true. Equally, ' $Raa$ ' is false: Aristotle was not born before Aristotle.

Dealing with atomic sentences, then, is very intuitive:

When  $R$  is an  $n$ -place predicate and  $a_1, a_2, \dots, a_n$  are names, the sentence  $Ra_1a_2\dots a_n$  is true in an interpretation iff  $R$  is true of the objects named by  $a_1, a_2, \dots, a_n$  (in that order) in that interpretation

Recall, though, that there is a second kind of atomic sentence: 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,

For any names  $a$  and  $b$ , the sentence  $a = b$  is true in an interpretation iff  $a$  and  $b$  name the very same object in that interpretation

So, in our go-to interpretation, ' $a = d$ ' is false, since ' $a$ ' names Aristotle and ' $d$ ' names Trump and Aristotle is distinct from Trump.

## 21.2 Sentential connectives

We saw in §19 that FOL sentences can be built up from simpler ones using the truth-functional connectives that were familiar from TFL. The rules governing these connectives are *exactly* the same as they were when we considered TFL. Here they are:

$A \wedge B$  is true in an interpretation iff  
both  $A$  is true and  $B$  is true in that interpretation

$A \vee B$  is true in an interpretation iff  
either  $A$  is true or  $B$  is true in that interpretation

$\neg A$  is true in an interpretation iff  
 $A$  is false in that interpretation

$A \rightarrow B$  is true in an interpretation iff  
either  $A$  is false or  $B$  is true in that interpretation

$A \leftrightarrow B$  is true in an interpretation iff  
 $A$  has the same truth value as  $B$  in that interpretation

This presents the very same information as the characteristic truth tables for the connectives; it just does it in a slightly different way. Some examples will probably help to illustrate the idea (make sure you understand them!) On our go-to interpretation:

- ' $a = a \wedge Wa$ ' is true
- ' $Rad \wedge Wd$ ' is false because, although ' $Rad$ ' is true, ' $Wd$ ' is false
- ' $a = d \vee Wa$ ' is true
- ' $\neg a = d$ ' is true
- ' $Wa \wedge \neg(a = d \wedge Rad)$ ' is true, because ' $Wa$ ' is true and ' $a = d$ ' is false

## 21.3 When the main logical operator is a quantifier

The exciting innovation in FOL, 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 (speaking roughly) ' $\exists yLxy$ ' is true of everything in the domain. But our interpretation does not *directly* specify what ' $\exists yLxy$ ' is true of. Instead, whether this is true of something should follow just from the interpretation of the predicate ' $L$ ', 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 our go-to interpretation only interprets *two* names, ‘ $a$ ’ and ‘ $d$ ’. But the domain—all people born before the year 2000CE—contains many more than two people. (And I have no intention of trying to correct for this by naming *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, ‘ $\exists xRdx$ ’ should be true. After all, in the domain, there is certainly someone who was born after Donald Trump. 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 ‘ $Rdc$ ’ would be true on this extended interpretation. And this, surely, should suffice to make ‘ $\exists xRdx$ ’ 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, ‘ $l$ ’, denote that person, then we would not be able to find anyone else in the domain to denote with some further new name, perhaps ‘ $m$ ’, in such a way that ‘ $Rlm$ ’ would be true. Indeed, no matter *whom* we named with ‘ $m$ ’, ‘ $Rlm$ ’ would be false. And this observation is surely sufficient to make ‘ $\exists yRly$ ’ *false* in our extended interpretation. And this is surely sufficient to make ‘ $\forall x\exists yRxy$ ’ false on the original go-to interpretation.

If you have understood these three examples, that’s what matters. It provides the basis for a formal definition of truth for quantified sentences.

Strictly speaking, though, we still need to *give* that definition. The result, sadly, is a bit ugly, and requires a few new definitions. Brace yourself!

Suppose that  $A$  is a formula containing at least one instance of the variable  $x$ , and that  $x$  is free in  $A$ . We will write this thus:

$$A(\dots x \dots x \dots)$$

Suppose also that  $c$  is a name. Then we shall write:

$$A(\dots c \dots c \dots)$$

for the formula we obtain by replacing *every* occurrence of  $x$  in  $A$  with  $c$ . The resulting formula is called a SUBSTITUTION INSTANCE of  $\forall xA$  and  $\exists xA$ . Also,  $c$  is 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'.

Armed with this notation, the rough idea is as follows. The sentence  $\forall \mathbf{x}A(\dots\mathbf{x}\dots\mathbf{x}\dots)$  will be true iff  $A(\dots\mathbf{c}\dots\mathbf{c}\dots)$  is true no matter what object (in the domain) we name with  $\mathbf{c}$ . Similarly, the sentence  $\exists \mathbf{x}A$  will be true iff there is *some* way to assign the name  $\mathbf{c}$  to an object (in the domain) that makes  $A(\dots\mathbf{c}\dots\mathbf{c}\dots)$  true. More precisely, we stipulate:

$\forall \mathbf{x}A(\dots\mathbf{x}\dots\mathbf{x}\dots)$  is true in an interpretation iff  $A(\dots\mathbf{c}\dots\mathbf{c}\dots)$  is true in *every* interpretation that extends the original interpretation by assigning an object to any name  $\mathbf{c}$  (without changing the interpretation in any other way).

$\exists \mathbf{x}A(\dots\mathbf{x}\dots\mathbf{x}\dots)$  is true in an interpretation iff  $A(\dots\mathbf{c}\dots\mathbf{c}\dots)$  is true in *some* interpretation that extends the original interpretation by assigning an object to some name  $\mathbf{c}$  (without changing the interpretation in any other way).

To be clear: all this is doing is formalising (very pedantically) the intuitive idea expressed on the previous page. The result is a bit ugly, and the final definition might look a bit opaque. Hopefully, though, the *spirit* of the idea is clear.

## Practice exercises

A. Consider the following interpretation:

- The domain comprises only Corwin and Benedict
- 'A' is a one-place predicate, to be true of both Corwin and Benedict
- 'B' is a one-place predicate, to be true of Benedict only
- 'N' is a one-place predicate, 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:

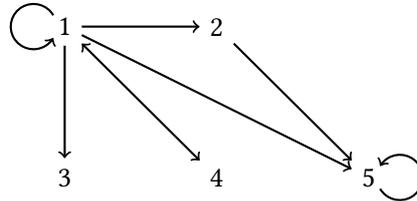
- The domain comprises only Lemmy, Courtney and Eddy
- 'G' is a one-place predicate, to be true of Lemmy, Courtney and Eddy.
- 'H' is a one-place predicate, to be true of and only of Courtney

- ‘ $M$ ’ is a one-place predicate, 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 §20, 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.  $\exists x \exists y (\neg x = y \wedge Rxy \wedge \forall z (Rzx \leftrightarrow y = z))$

# Semantic concepts

22

Defining truth in FOL was quite fiddly. But now that we are done, we can define various other central logical notions. These definitions will look very similar to those for TFL, from §11. However, remember that they concern *interpretations*, rather than valuations.

We will use the symbol ‘ $\models$ ’ for FOL much as we did for TFL. So:

$$A_1, A_2, \dots, A_n \models C$$

means that no interpretation makes all of  $A_1, A_2, \dots, A_n$  true and  $C$  is false.

Derivatively, here is how to say that  $C$  is true in every interpretation:

$$\models C$$

In this case, we say that  $C$  is a LOGICAL TRUTH.

Equally, here is how to say that  $A$  is false in every interpretation:

$$A \models$$

In this case, we say that  $A$  is a CONTRADICTION.

To say that it is *not* the case that  $A_1, \dots, A_n \models C$ , we slash the turnstile and write:

$$A_1, A_2, \dots, A_n \not\models C$$

This means that some interpretation makes all of  $A_1, \dots, A_n$  true whilst making  $C$  false.

Two FOL sentences  $A$  and  $B$  are LOGICALLY EQUIVALENT iff they are true in exactly the same interpretations as each other; i.e., both  $A \models B$  and  $B \models A$ .

Some FOL sentences are JOINTLY CONSISTENT iff some interpretation makes all of them true. They are JOINTLY INCONSISTENT iff there is no such interpretation.

# Using interpretations

## 23.1 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 *just* 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 just offer:

$A$ : \_\_\_\_<sub>1</sub> is identical with \_\_\_\_<sub>2</sub>

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:

$B$ : \_\_\_\_<sub>1</sub> 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

$A$ : \_\_\_\_<sub>1</sub> is identical with \_\_\_\_<sub>2</sub>

$B$ : \_\_\_\_<sub>1</sub> 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.

## 23.2 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. *Exercise*: work out why!) For concreteness, let’s take:

domain: Ornette Coleman, Miles Davis

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, let’s say:

$S$ : \_\_\_\_<sub>1</sub> plays saxophone

Now ‘ $\exists xSx$ ’ is true, because ‘ $S$ ’ 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 ‘ $S$ ’ is false of Miles Davis. Slightly more precisely, extend our interpretation by allowing ‘ $d$ ’ to name Miles Davis, 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  $A$  is not a logical truth, it suffices to find an interpretation where  $A$  is false.

To show that  $A$  is not a contradiction, it suffices to find an interpretation where  $A$  is true.

To show that  $A$  and  $B$  are not logically equivalent, it suffices to find an interpretation where one is true and the other is false.

## 23.3 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 FOL:

$$\exists x(Gx \rightarrow Ga) \therefore \exists xGx \rightarrow Ga$$

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

$G$ : \_\_\_\_<sub>1</sub> 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 ‘ $\exists xGx \rightarrow Ga$ ’, i.e. that  $\exists x(Gx \rightarrow Ga) \not\equiv \exists xGx \rightarrow Ga$ . 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:

$$\forall x\exists yLxy \therefore \exists y\forall xLxy$$

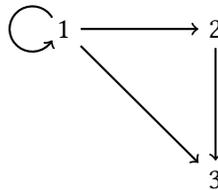
Again, I want to show that this is invalid. To do this, we must make the premises true and the conclusion false. Here is a suggestion:

domain: UK citizens currently in a civil partnership with another UK citizen

$L$ : \_\_\_\_\_<sub>1</sub> is in a civil partnership with \_\_\_\_\_<sub>2</sub>

The premise is clearly true on this interpretation. Anyone in the domain is a UK citizen in a civil partnership with some other UK citizen. That other citizen will also, then, be in the domain. So for everyone in the domain, there will be someone (else) in the domain with whom they are in a civil partnership. Hence ‘ $\forall x\exists yLxy$ ’ is true. But the conclusion is clearly false, for that would require that there is some single person who is in a civil partnership 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 §20, I described how we can present some interpretations using diagrams. For example:



Using the conventions employed in §20, 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.

If some interpretation makes all of  $A_1, A_2, \dots, A_n$  true and  $C$  is false, then:

- $A_1, A_2, \dots, A_n \therefore C$  is *invalid*; and
- $A_1, A_2, \dots, A_n \not\models C$ ; and
- And  $A_1, A_2, \dots, A_n, \neg C$  are jointly consistent.

An interpretation which refutes a claim—to logical truth, say, or to entailment—is called a *counter-interpretation*, or a *counter-model*.

I should close this section, though, with a caution about the relationship between (in)validity and (non)entailment. Recall FOL's limitations: it is an extensional language; it ignores issues of vagueness; and it cannot handle cases of validity for 'special reasons'. To take one illustration of these issues, consider this natural-language argument:

Every vixen will die

So: All foxes will die

This is valid: necessarily every vixen is a fox, so it is impossible for the premise to be true and the conclusion false. Now, we might sensibly symbolise the argument as follows:

$$\forall x(Vx \rightarrow Dx) \therefore \forall x(Fx \rightarrow Dx)$$

However, it is easy to find counter-models which show that  $\forall x(Vx \rightarrow Dx) \not\models \forall x(Fx \rightarrow Dx)$ . (*Exercise*: find one.) So, it would be *wrong* to infer that the English argument is *invalid*, just because there is a counter-model to the relevant *FOL-entailment*.

The general moral is this. If you want to infer from the absence of an entailment in FOL to the invalidity of some English argument, then you need to argue that nothing important is lost in the way you have symbolised the English argument.

## 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.  $\forall z\exists yz \leftrightarrow \exists y\forall yz$
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.  $\exists a, Ka$
2.  $\exists x\exists x, \exists m$
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.  $\forall x\exists yRxy, \exists x\forall yRxy$

$$9. \forall x \exists y Rxy, \forall x \exists y Ryx$$

C. Show that the following sentences are jointly consistent:

1.  $Ma, \neg Na, Pa, \neg Qa$
2.  $Lee, Leg, \neg Lge, \neg Lgg$
3.  $\neg(Ma \wedge \exists x Ax), Ma \vee Fa, \forall x(Fx \rightarrow Ax)$
4.  $Ma \vee Mb, Ma \rightarrow \forall x \neg Mx$
5.  $\forall y Gy, \forall x(Gx \rightarrow Hx), \exists y \neg Iy$
6.  $\exists x(Bx \vee Ax), \forall x \neg Cx, \forall x[(Ax \wedge Bx) \rightarrow Cx]$
7.  $\exists x Xx, \exists x Yx, \forall x(Xx \leftrightarrow \neg Yx)$
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.  $\neg \exists x \forall y Rxy, \forall x \exists y Rxy$
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.  $\forall x(Ax \rightarrow Bx) \therefore \exists x Bx$
2.  $\forall x(Rx \rightarrow Dx), \forall x(Rx \rightarrow Fx) \therefore \exists x(Dx \wedge Fx)$
3.  $\exists x(Px \rightarrow Qx) \therefore \exists x Px$
4.  $Na \wedge Nb \wedge Nc \therefore \forall x Nx$
5.  $Rde, \exists x Rxd \therefore Red$
6.  $\exists x(Ex \wedge Fx), \exists x Fx \rightarrow \exists x Gx \therefore \exists x(Ex \wedge Gx)$
7.  $\forall x Oxc, \forall x Ocx \therefore \forall x Oxx$
8.  $\exists x(\neg jx \wedge Kx), \exists x \neg Kx, \exists x \neg jx \therefore \exists x(\neg jx \wedge \neg Kx)$
9.  $Lab \rightarrow \forall x Lxb, \exists x Lxb \therefore Lbb$
10.  $\forall x(Dx \rightarrow \exists y Tyx) \therefore \exists y \exists z \neg y = z$

# Reasoning about all interpretations 24

## 24.1 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 FOL. Rather, it is an argument in English *about* FOL: 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 TFL sentences.)

Let's have another example. The sentence ' $\forall x(Rxx \leftrightarrow Rxx)$ ' should obviously be a logical truth. However, saying *precisely* why is quite tricky. We cannot say that ' $Rxx \leftrightarrow Rxx$ ' is true in every interpretation, since ' $Rxx \leftrightarrow Rxx$ ' is not even a *sentence* of FOL (remember that ' $x$ ' is a variable, not a name). Instead, we should say something like this.

Consider some arbitrary interpretation. Consider some arbitrary member of the interpretation's domain—for convenience, call it *obbie*—and suppose that 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.

## 24.2 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! For a final example, here is a perfectly obvious entailment:

$$\forall x(Hx \wedge Jx) \models \forall xHx$$

After all, if everything is both H and J, then everything is H. But we can only establish the entailment 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.

Even for a simple entailment like this one, the reasoning is somewhat complicated. For more complicated entailments, the reasoning can be extremely torturous.

The following table summarises whether a single (counter-)interpretation suffices, or whether we must reason about all interpretations.

	Yes	No
logical truth?	all interpretations	one counter-interpretation
contradiction?	all interpretations	one counter-interpretation
equivalent?	all interpretations	one counter-interpretation
consistent?	one interpretation	all interpretations
valid?	all interpretations	one counter-interpretations
entailment?	all interpretations	one counter-interpretation

---

You might want to compare this table with the table at the end of §13. The key difference resides in the fact that TFL concerns truth tables, whereas FOL 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.

6

## Natural deduction for TFL

# The very idea of natural deduction 25

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 TFL, this led us to develop truth tables. Each line of a complete truth table corresponds to a valuation. So, when faced with a TFL 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 may not give us much *insight*. Consider two arguments in TFL:

$$\begin{aligned} P \vee Q, \neg P \therefore Q \\ P \rightarrow Q, P \therefore Q \end{aligned}$$

Clearly, these are valid arguments. You can confirm that they are valid by constructing four-line truth tables. But we might say that they make use of different *forms* of reasoning. And it might be nice to keep track of these different forms of inference.

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. We begin with very basic rules of inference. These rules can be combined, to offer more complicated arguments. Indeed, with just a small starter pack of rules of inference, 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.

The move to natural deduction might be motivated by more than the search for insight. It might also be motivated by *necessity*. 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 TFL can become *intractable* if we use truth tables.

When we get to FOL, though, the problem gets dramatically worse. There is nothing like the truth table test for FOL. 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.

There are, indeed, systems that codify ways to reason about all possible interpretations. They were developed in the 1950s by Evert Beth and Jaakko Hintikka. But we shall not follow this path. We shall, instead, look to natural deduction.

Rather than reasoning directly about all valuations (in the case of TFL) or all interpretations (in the case of FOL), we shall try to select a few basic rules of inference. 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).

I 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 logical operator, and ELIMINATION rules, that allow us to prove something given a sentence that has that connective as the main logical operator

## 26.1 The idea of a formal proof

A *formal proof* is a sequence of sentences, some of which are marked as being initial assumptions (or premises). The last line of the formal proof is the conclusion. (Henceforth, I shall simply call these ‘proofs’, but be aware that there are *informal proofs* too.)

As an illustration, consider:

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

We shall start a proof by writing the premise:

$$1 \quad \underline{\neg(A \vee B)}$$

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. 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 what line number 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:

$$\begin{array}{l} 1 \quad A \vee B \\ 2 \quad \neg(A \wedge C) \\ 3 \quad \underline{\neg(B \wedge \neg D)} \end{array}$$

and we are hoping to conclude with some line:

$$n \quad | \quad \neg C \vee D$$

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.

## 26.2 Conjunction

Suppose I want to show that Jacob is both hypocritical and reactionary. One obvious way to do this would be as follows: first I show that Jacob is hypocritical; then I show that Jacob is reactionay; 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:

*H*: Jacob is hypocritical

*R*: Jacob is reactionary

Perhaps I am working through a proof, and I have obtained '*H*' on line 8 and '*R*' on line 15. Then on any subsequent line I can obtain '*H* ∧ *R*' thus:

$$\begin{array}{l|l} 8 & H \\ 15 & R \\ & H \wedge R \quad \wedge I 8, 15 \end{array}$$

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. I could equally well obtain:

$$\begin{array}{l|l} 8 & H \\ 15 & R \\ & R \wedge H \quad \wedge I 15, 8 \end{array}$$

with the citation reversed, to reflect the order of the conjuncts. More generally, here is our conjunction introduction rule:

$\begin{array}{l l} m & \mathbf{A} \\ n & \mathbf{B} \\ & \mathbf{A} \wedge \mathbf{B} \quad \wedge I m, n \end{array}$
---

To be clear, the statement of the rule is *schematic*. It is not itself a proof. '*A*' and '*B*' are not sentences of TFL. Rather, they are symbols in the metalanguage, which we use when we want to talk about any sentence of TFL (see §7). Similarly, '*m*' and '*n*' are not a 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. They can be separated from one another, and they can appear in any order.

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 Jacob is both hypocritical and reactionary. Then you are entitled to conclude that Jacob is hypocritical. Equally, you are entitled to conclude that Jacob is reactionary. Putting this together, we obtain our conjunction elimination rule(s):

$$\begin{array}{l|l} m & A \wedge B \\ & A \quad \wedge E\ m \end{array}$$

and equally:

$$\begin{array}{l|l} m & A \wedge B \\ & B \quad \wedge E\ m \end{array}$$

The point is simply that, when you have a conjunction on some line of a proof, you can obtain either of the conjuncts by  $\wedge E$ . (But one point is worth emphasising: you can only apply this rule when conjunction is the main logical operator. So you cannot infer ' $D$ ' 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:

$$\begin{array}{l} [(A \vee B) \rightarrow (C \vee D)] \wedge [(E \vee F) \rightarrow (G \vee H)] \\ \therefore [(E \vee F) \rightarrow (G \vee H)] \wedge [(A \vee B) \rightarrow (C \vee D)] \end{array}$$

The main logical operator 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 (repeated applications of) our rules of inference. 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:

1	$[(A \vee B) \rightarrow (C \vee D)] \wedge [(E \vee F) \rightarrow (G \vee H)]$	
2	$[(A \vee B) \rightarrow (C \vee D)]$	$\wedge E$ 1
3	$[(E \vee F) \rightarrow (G \vee H)]$	$\wedge E$ 1
4	$[(E \vee F) \rightarrow (G \vee H)] \wedge [(A \vee B) \rightarrow (C \vee D)]$	$\wedge I$ 3, 2

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 256 lines; our formal proof required only four lines.

It is worth giving another example. 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 with this argument, we start by writing:

1	$A \wedge (B \wedge C)$
---	-------------------------

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:

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

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 TFL 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.)

Let me offer one final illustration. When using the  $\wedge$ I rule, there is no need to apply it to different sentences. So, if we want, we can formally prove ‘ $A$ ’ from ‘ $A$ ’ thus:

1	$A$	
2	$A \wedge A$	$\wedge$ I 1, 1
3	$A$	$\wedge$ E 2

### 26.3 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):

$m$	$A \rightarrow B$	
$n$	$A$	
	$B$	$\rightarrow$ 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:

Jacob is hypocritical. So: if Jacob is reactionary, then Jacob is both hypocritical and reactionary.

If someone doubted that this was valid, we might try to convince them otherwise by explaining ourselves as follows:

Assume that Jacob is hypocritical. Now, *additionally* assume that Jacob is reactionary. Then by conjunction introduction—which we just discussed—Jacob is both hypocritical and reactionary. Of course, that’s conditional on the additional assumption that Jacob is reactionary. But this just means that, if Jacob is reactionary, then he is both hypocritical and reactionary.

Transferred into natural deduction format, here is the pattern of reasoning that we just used. We started with one premise, ‘Jacob is hypocritical’, thus:

$$1 \quad \underline{H}$$

The next thing we did is to make an *additional* assumption ('Jacob is reactionary'), for the sake of argument. To indicate that we are no longer dealing *merely* with our original assumption (' $H$ '), but with some additional assumption, we continue our proof as follows:

$$\begin{array}{l} 1 \quad \underline{H} \\ 2 \quad \underline{\quad} \underline{R} \end{array}$$

Note that we are *not* claiming, on line 2, to have proved ' $R$ ' 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:

$$\begin{array}{l} 1 \quad \underline{H} \\ 2 \quad \underline{\quad} \underline{R} \\ 3 \quad \underline{\quad} \underline{H \wedge R} \quad \wedge I 1, 2 \end{array}$$

So we have now shown that, on the additional assumption, ' $R$ ', we can obtain ' $H \wedge R$ '. We can therefore conclude that, if ' $R$ ' obtains, then so does ' $H \wedge R$ '. Or, to put it more briefly, we can conclude ' $R \rightarrow (H \wedge R)$ ':

$$\begin{array}{l} 1 \quad \underline{H} \\ 2 \quad \underline{\quad} \underline{R} \\ 3 \quad \underline{\quad} \underline{H \wedge R} \quad \wedge I 1, 2 \\ 4 \quad R \rightarrow (H \wedge R) \quad \rightarrow I 2-3 \end{array}$$

Observe that we have dropped back to using one vertical line. We have *discharged* the additional assumption, ' $R$ ', since the conditional itself follows just from our original assumption, ' $H$ '.

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 know the following: If  $A$ , then  $B$ . This is wrapped up in the rule for conditional introduction:

$\begin{array}{l} i \quad \underline{\quad} \underline{\quad} \underline{A} \\ j \quad \underline{\quad} \underline{\quad} \underline{B} \\ \underline{\quad} \underline{A \rightarrow B} \quad \rightarrow I i-j \end{array}$
--

There can be as many or as few lines as you like between lines  $i$  and  $j$ .

It will help to offer a second 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:

$$\begin{array}{l|l} 1 & P \rightarrow Q \\ 2 & Q \rightarrow R \\ \hline 3 & \begin{array}{|l} P \end{array} \end{array}$$

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:

$$\begin{array}{l|l|l} 1 & P \rightarrow Q & \\ 2 & Q \rightarrow R & \\ \hline 3 & \begin{array}{|l} P \end{array} & \\ 4 & \begin{array}{|l} Q \end{array} & \rightarrow E 1, 3 \\ 5 & \begin{array}{|l} R \end{array} & \rightarrow E 2, 4 \\ 6 & P \rightarrow R & \rightarrow I 3-5 \end{array}$$

## 26.4 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:

$$\begin{array}{l|l|l} 1 & A & \\ \hline 2 & \begin{array}{|l} B \end{array} & \\ 3 & \begin{array}{|l} B \wedge B \end{array} & \wedge I 2, 2 \\ 4 & \begin{array}{|l} B \end{array} & \wedge E 3 \\ 5 & B \rightarrow B & \rightarrow I 2-4 \end{array}$$

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 ∧ B	∧I 2, 2
4	B	∧E 3
5	B → B	→I 2–4
6	B	naughty attempt to invoke →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-foot tall (symbolised by ‘B’)! So we must be prohibited from doing this. 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:

To cite individual lines when applying a rule, those lines must (1) come before the application of the rule, but (2) not occur within a closed subproof.

This stipulation rules out the disastrous attempted proof above. The rule of →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 using only the rules which 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 infer 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:

To cite a subproof when applying a rule, the subproof must (1) come before the application of the rule, but (2) not occur within some other closed subproof.

The attempted disastrous proof violates this stipulation. The subproof of lines 3–4 occurs within the subproof of lines 2–5, so the subproof of lines 3–4 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 →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.

## 26.5 Biconditional

The rules for the biconditional will be like double-barrelled versions of the rules for the conditional.

In order to prove ' $F \leftrightarrow G$ ', for instance, you must be able to prove ' $G$ ' on the assumption ' $F$ ' and prove ' $F$ ' on the assumption ' $G$ '. The biconditional introduction rule ( $\leftrightarrow$ I) therefore requires two subproofs. Schematically, the rule works like this:

$i$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>A</math></td> </tr> </table>	$A$	
$A$			
$j$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>B</math></td> </tr> </table>	$B$	
$B$			
$k$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>B</math></td> </tr> </table>	$B$	
$B$			
$l$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>A</math></td> </tr> </table>	$A$	
$A$			
	$A \leftrightarrow B$	$\leftrightarrow$ 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.

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 allow:

$m$	$A \leftrightarrow B$	
$n$	$A$	
	$B$	$\leftrightarrow$ E $m, n$

and equally:

$m$	$A \leftrightarrow B$	
$n$	$B$	
	$A$	$\leftrightarrow$ 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.

## 26.6 Disjunction

Suppose Jacob is hypocritical. Then certainly Jacob is either hypocritical or reactionary. After all, to say that Jacob is either hypocritical or reactionary is to say something weaker

than to say that Jacob is hypocritical.

Let me emphasise this point. Suppose Jacob is hypocritical. It follows that Jacob is *either* hypocritical *or* a kumquat. Equally, it follows that *either* Jacob is hypocritical *or* that kumquats are the only fruit. Equally, it follows that *either* Jacob is hypocritical *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 violate lots of implicit conversational norms).

Armed with all this, I present the disjunction introduction rule(s):

$m$	$A$		
	$A \vee B$		$\vee I m$

and

$m$	$A$		
	$B \vee A$		$\vee I m$

Notice that  $B$  can be *any* sentence whatsoever. So the following is a perfectly kosher proof:

1		$M$		
		$M \vee ((A \leftrightarrow B) \rightarrow (C \wedge D)) \leftrightarrow [E \wedge F]$		$\vee I 1$

Using a truth table to show this would have taken 128 lines.

The disjunction elimination rule is, though, slightly trickier. Suppose that either Jacob is hypocritical or he is reactionary. What can you conclude? Not that Jacob is hypocritical; it might be that he is reactionary instead. And equally, not that Jacob is reactionary; for he might merely be hypocritical. Disjunctions, just by themselves, are hard to work with.

But suppose that we could somehow show both of the following: first, that Jacob's being hypocritical entails that he is a brexiteer; second, that Jacob's being reactionary entails that he is a brexiteer. Then if we know that Jacob is either hypocritical or reactionary, then we know that, whichever he happens to be, Jacob is a brexiteer. This insight can be expressed in the following rule, which is our disjunction elimination ( $\vee E$ ) rule:

$m$	$A \vee B$		
$i$		$A$	
$j$		$C$	
$k$		$B$	
$l$		$C$	
	$C$		$\vee E m, i-j, k-l$

This is obviously a bit clunkier to write down than our previous rules, but the point is fairly simple. Suppose we have some disjunction,  $A \vee B$ . Suppose we have two subproofs, showing us that  $C$  follows from the assumption that  $A$ , and that  $C$  follows from the assumption that  $B$ . Then we can infer  $C$  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: 10px;"> <math>P \wedge Q</math> </td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>P</math> </td> <td style="padding-left: 20px;"><math>\wedge E</math> 2</td> </tr> </table>	$P \wedge Q$		$P$	$\wedge E$ 2	
$P \wedge Q$						
$P$	$\wedge E$ 2					
3	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>P \wedge R</math> </td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>P</math> </td> <td style="padding-left: 20px;"><math>\wedge E</math> 4</td> </tr> </table>	$P \wedge R$		$P$	$\wedge E$ 4	
$P \wedge R$						
$P$	$\wedge E$ 4					
4	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>P</math> </td> <td style="padding-left: 20px;"><math>\vee E</math> 1, 2-3, 4-5</td> </tr> </table>	$P$	$\vee E$ 1, 2-3, 4-5			
$P$	$\vee E$ 1, 2-3, 4-5					
5	$P$					
6						

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	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>B</math> </td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>A \wedge B</math> </td> <td style="padding-left: 20px;"><math>\wedge I</math> 2, 4</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>(A \wedge B) \vee (A \wedge C)</math> </td> <td style="padding-left: 20px;"><math>\vee I</math> 5</td> </tr> </table>	$B$		$A \wedge B$	$\wedge I$ 2, 4	$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 5	
$B$								
$A \wedge B$	$\wedge I$ 2, 4							
$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 5							
5	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>C</math> </td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>A \wedge C</math> </td> <td style="padding-left: 20px;"><math>\wedge I</math> 2, 7</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"> <math>(A \wedge B) \vee (A \wedge C)</math> </td> <td style="padding-left: 20px;"><math>\vee I</math> 8</td> </tr> </table>	$C$		$A \wedge C$	$\wedge I$ 2, 7	$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 8	
$C$								
$A \wedge C$	$\wedge I$ 2, 7							
$(A \wedge B) \vee (A \wedge C)$	$\vee I$ 8							
6	$(A \wedge B) \vee (A \wedge C)$	$\vee E$ 3, 4-6, 7-9						
7								
8								
9								
10								

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 comes 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.

## 26.7 Contradiction and negation

We have only one connective left to deal with: negation. But to tackle it, I must connect negation with *contradiction*.

An effective form of argument is to argue your opponent into contradicting themselves. At that point, you have them on the ropes. They have to give up at least one of their assumptions. We are going to make use of this idea in our proof system, by adding a new symbol, ‘ $\perp$ ’, to our proofs. This should be read as something like ‘contradiction!’ or ‘reductio!’ or ‘but that’s absurd!’ And the rule for introducing this symbol is that we can use it whenever we explicitly contradict ourselves, i.e. whenever we find both a sentence and its negation appearing in our proof:

$m$	$A$	
$n$	$\neg A$	
	$\perp$	$\neg E\ m, n$

It does not matter what order the sentence and its negation appear in, and they do not need to appear on adjacent lines. However, we always cite the sentence first, followed by its negation. The rule here is called  $\neg E$ , since the negation sign is eliminated. (We could equally have called this rule ‘ $\perp I$ ’, since it introduces ‘ $\perp$ ’.)

Next, we need a rule for negation introduction. The rule is very simple: if assuming something leads you to a contradiction, then the assumption must be wrong. This thought motivates the following rule:

$i$	<table style="border-left: 1px solid black; padding-left: 10px;"> <tr> <td style="padding-right: 10px;"><math>A</math></td> </tr> <tr> <td style="border-top: 1px solid black; padding-top: 5px;"><math>\perp</math></td> </tr> </table>	$A$	$\perp$	
$A$				
$\perp$				
$j$	$\neg A$	$\neg I\ i-j$		

There can be as many lines between  $i$  and  $j$  as you like. To see our negation introduction and elimination rules in practice, consider this proof:

1	$D$			
2	<table style="border-left: 1px solid black; padding-left: 10px;"> <tr> <td style="padding-right: 10px;"><math>\neg D</math></td> </tr> <tr> <td style="border-top: 1px solid black; padding-top: 5px;"><math>\perp</math></td> </tr> </table>	$\neg D$	$\perp$	
$\neg D$				
$\perp$				
3	$\perp$	$\neg E\ 1, 2$		
4	$\neg\neg D$	$\neg I\ 2-3$		

So far so good. However, now that we have a symbol for contradiction, ‘ $\perp$ ’, we need a rule for eliminating it. The rule we will use is *ex falso quod libet*. This means ‘anything follows from a contradiction’. And the idea is precisely that: if we obtained a contradiction, symbolised by ‘ $\perp$ ’, then we can infer whatever we like. For just this reason, the rule is sometimes also called *explosion*: in the presence of a contradiction, inference explodes.

How can this be motivated, informally? Well, consider the English rhetorical device ‘...and if *that’s* true, I’ll eat my hat’. Contradictions simply cannot be true. So, if one is true, then not only will I eat my hat, I’ll have it too.<sup>1</sup> Here is the formal rule:

$m$	$\perp$	$A$	$X m$
-----	---------	-----	-------

Think of the ‘X’ as standing for ‘eX falso’, or ‘eXplosion’, and note that  $A$  can be *any* sentence whatsoever.

Now, I have said that ‘ $\perp$ ’ should be read as something like ‘contradiction!’ But this does not tell us much about the symbol. There are, roughly, three ways to approach the symbol.

- We might regard ‘ $\perp$ ’ as a new atomic sentence of TFL, but one which can only ever have the truth value False.
- We might regard ‘ $\perp$ ’ as an abbreviation for some canonical contradiction, such as ‘ $A \wedge \neg A$ ’. This will have the same effect as the above—obviously, ‘ $A \wedge \neg A$ ’ only ever has the truth value False—but it means that, officially, we do not need to add a new symbol to TFL.
- We might regard ‘ $\perp$ ’, not as a symbol of TFL, but as something more like a *punctuation mark* that appears in our proofs. (It is on a par with the line numbers and the vertical lines, say.)

There is something very philosophically attractive about the third option. But here I shall *officially* plump for the second. ‘ $\perp$ ’ is to be read as abbreviating some canonical contradiction. This means that we can manipulate it, in our proofs, just like any other sentence.

## 26.8 Tertium non datur

I will add one final rule, governing negation. It is very much like the rule used in disjunction elimination, and it requires a little motivation.

Suppose that we can show that if it’s sunny outside, then Bill will have brought an umbrella (for fear of sunbearn). 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:

$i$	$A$	
$j$	$B$	
$k$	$\neg A$	
$l$	$B$	
	$B$	$\text{TND } i-j, k-l$

<sup>1</sup>Thanks to Adam Caulton for this.



1	$P \wedge S$
2	$S \rightarrow R$
3	$P$
4	$S$
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	$\perp$
9	$J$
10	$J$

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.  $\neg(P \rightarrow Q) \therefore P$

# Additional rules for TFL

In §26, we introduced the basic rules of our proof system for TFL. In this section, we shall add some additional rules to our system. These will make our system much easier to work with. (However, in §30 we will see that they are not strictly speaking *necessary*.)

## 27.1 Reiteration

The first additional rule is *reiteration* (R). This just allows us to repeat ourselves:

$\begin{array}{l l} m & A \\ & A \quad R\ m \end{array}$
--

Such a rule is obviously legitimate; but one might well wonder how such a rule could ever be useful. Well, consider this:

1	$F \vee G$	
2	$F \rightarrow G$	
3	$F$	
4	$G$	$\rightarrow E\ 2, 3$
5	$G$	
6	$G$	$R\ 5$
7	$G$	$\vee E\ 1, 3-4, 5-6$

## 27.2 Disjunctive syllogism

Here is a very natural argument form.

Trump is either in DC or Mar-a-Lago. He is not in DC. So, he is in Mar-a-Lago.

This is called *disjunctive syllogism*. We add it to our proof system as follows:

$m$	$A \vee B$	
$n$	$\neg A$	
	$B$	DS $m, n$

and

$m$	$A \vee B$	
$n$	$\neg B$	
	$A$	DS $m, n$

As usual, the disjunction and the negation of one disjunct may occur in either order and need not be adjacent. However, we always cite the disjunction first. (This is, if you like, a new rule of disjunction elimination.)

### 27.3 Modus tollens

Another useful pattern of inference is embodied in the following argument:

If Hillary has won the election, then she is in the White House. She is not in the White House. So she has not won the election.

This inference pattern is called *modus tollens*. The corresponding rule is:

$m$	$A \rightarrow B$	
$n$	$\neg B$	
	$\neg A$	MT $m, n$

As usual, the premises may occur in either order, but we always cite the conditional first. (This is, if you like, a new rule of conditional elimination.)

### 27.4 Double-negation elimination

Another useful rule is *double-negation elimination*. It does exactly what it says on the tin:

$m$	$\neg\neg A$	
	$A$	DNE $m$

The justification for this is that, in natural language, 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 infer '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 TFL forces us to sacrifice certain nuances of English expressions.

## 27.5 De Morgan Rules

Our final additional rules are called De Morgan's Laws (named after August De Morgan). The shape of the rules should be familiar from truth tables.

The first De Morgan rule is:

$$\begin{array}{l|l} m & \neg(A \wedge B) \\ & \neg A \vee \neg B \quad \text{DeM } m \end{array}$$

The second De Morgan is the reverse of the first:

$$\begin{array}{l|l} m & \neg A \vee \neg B \\ & \neg(A \wedge B) \quad \text{DeM } m \end{array}$$

The third De Morgan rule is the *dual* of the first:

$$\begin{array}{l|l} m & \neg(A \vee B) \\ & \neg A \wedge \neg B \quad \text{DeM } m \end{array}$$

And the fourth is the reverse of the third:

$$\begin{array}{l|l} m & \neg A \wedge \neg B \\ & \neg(A \vee B) \quad \text{DeM } m \end{array}$$

*These are all of the additional rules of our proof system for TFL.*

## Practice exercises

A. The following proofs are missing their citations (rule and line numbers). Add them wherever they are required:

1	$W \rightarrow \neg B$
2	$A \wedge W$
3	$B \vee (J \wedge K)$
4	$W$
5	$\neg B$
6	$J \wedge K$
7	$K$

1	$L \leftrightarrow \neg O$
2	$L \vee \neg O$
3	$\neg L$
4	$\neg O$
5	$L$
6	$\perp$
7	$\neg\neg L$
8	$L$

1	$Z \rightarrow (C \wedge \neg N)$
2	$\neg Z \rightarrow (N \wedge \neg C)$
3	$\neg(N \vee C)$
4	$\neg N \wedge \neg C$
5	$\neg N$
6	$\neg C$
7	$Z$
8	$C \wedge \neg N$
9	$C$
10	$\perp$
11	$\neg Z$
12	$N \wedge \neg C$
13	$N$
14	$\perp$
15	$\neg\neg(N \vee C)$
16	$N \vee C$

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.  $(X \wedge Y) \vee (X \wedge Z), \neg(X \wedge D), D \vee M \therefore M$

# Proof-theoretic concepts

I now want to introduce some new vocabulary. The following expression:

$$A_1, A_2, \dots, A_n \vdash C$$

means that there is some proof which ends with  $C$  whose undischarged assumptions are among  $A_1, A_2, \dots, A_n$ . When we want to say that it is *not* the case that there is some proof which ends with  $C$  from  $A_1, A_2, \dots, A_n$ , we write:

$$A_1, A_2, \dots, A_n \nvdash C$$

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 FOL). *They are different notions.*

Armed with our new symbol, ‘ $\vdash$ ’, we can introduce some more terminology. To say that there is a proof of  $A$  with no undischarged assumptions, we write:

$$\vdash A$$

In this case, we say that  $A$  is a **THEOREM**.

To illustrate this, suppose I want to show that  $\vdash \neg(A \wedge \neg A)$ . So I need a proof of ‘ $\neg(A \wedge \neg A)$ ’ which has *no* undischarged assumptions. However, since I want to prove a sentence whose main logical operator is a negation, I shall want to start with a *subproof* within which I assume ‘ $A \wedge \neg A$ ’ and show that this assumption leads to contradiction. 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	$\perp$	$\neg E$ 2, 3
5	$\neg(A \wedge \neg A)$	$\neg I$ 1-4

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*.

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  $A$  and  $B$  are PROVABLY EQUIVALENT iff each can be proved from the other; i.e., both  $A \vdash B$  and  $B \vdash A$ .

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.

Here is a third, related, bit of terminology:

The sentences  $A_1, A_2, \dots, A_n$  are JOINTLY CONTRARY iff a contradiction can be proved from them, i.e.  $A_1, A_2, \dots, A_n \vdash \perp$ .

It is easy to show that some sentences are jointly contrary: you just need to prove a 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
contrary?	one proof	all possible proofs

## 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  $A \vdash B$ , what can you say about  $(A \wedge C) \vdash B$ ? What about  $(A \vee C) \vdash B$ ? 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  $A$  and  $B$  were provably equivalent.)

# Proof strategies

29

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 logical operator. This gives you an idea of what should happen *just before* the last line of the proof. Then you can treat this line as if it were your goal. And you can ask what you could do to get to this new goal.

For example: If your conclusion is a conditional  $A \rightarrow B$ , plan to use the  $\rightarrow$ I rule. This requires starting a subproof in which you assume  $A$ . The subproof ought to end with  $B$ . So, what can you do to get  $B$ ?

**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  $A$  directly, try starting by assuming  $\neg A$ . If a contradiction follows, then you will be able to obtain  $\neg\neg A$  by  $\neg$ I, and then  $A$  by DNE.

**Persist.** Try different things. If one approach fails, then try something else.

In this section, I will explain why I introduced the rules of our proof system in two separate batches. In particular, I want to show that the additional rules of §27 are not strictly speaking necessary, but can be derived from the basic rules of §26.

## 30.1 Derivation of Reiteration

Suppose you have some sentence on some line of your deduction:

$$m \quad | \quad A$$

You now want to repeat yourself, on some line  $k$ . You could just invoke the rule R, introduced in §27. But equally well, you can do this with the *basic* rules of §26:

$$\begin{array}{l|l} m & A \\ k & A \wedge A \quad \wedge I \ m \\ k+1 & A \quad \wedge E \ k \end{array}$$

To be clear: this is not a proof. Rather, it is a proof *scheme*. After all, it uses a variable, ‘ $A$ ’, rather than a sentence of TFL. But the point is simple. Whatever sentences of TFL we plugged in for ‘ $A$ ’, and whatever lines we were working on, we could produce a bona fide proof. So you can think of this as a recipe for producing proofs.

Indeed, it is a recipe which shows us the following: anything we can prove using the rule R, we can prove (with one more line) using just the *basic* rules of §26. So we can describe rule R as a *derived* rule, since it can be justified using only the basic rules.

## 30.2 Derivation of Disjunctive syllogism

Suppose that you are in a proof, and you have something of this form:

$$\begin{array}{l|l} m & A \vee B \\ n & \neg A \end{array}$$

You now want, on line  $k$ , to prove  $B$ . You can do this with the rule of DS, introduced in §27. But equally well, you can do this with the *basic* rules of §26:

$m$	$A \vee B$	
$n$	$\neg A$	
$k$	$A$	
$k+1$	$\perp$	$\neg E\ k, n$
$k+2$	$B$	$X\ k+1$
$k+3$	$B$	
$k+4$	$B \wedge B$	$\wedge I\ k+3, k+3$
$k+5$	$B$	$\wedge E\ k+4$
$k+6$	$B$	$\vee E\ m, k-k+2, k+3-k+5$

So the DS rule, again, can be derived from our more basic rules. Adding it to our system did not make any new proofs possible. Anytime you use the DS rule, you could always prove the same thing using a few extra lines but only the basic rules. As such, DS is a *derived* rule.

Similar comments will apply to all of the other rules flagged as ‘derived rules’. Indeed, I just need to work through each of modus tollens, the De Morgan rules, and double-negation elimination.

### 30.3 Derivation of Modus tollens

Suppose you have the following in your proof:

$m$	$A \rightarrow B$
$n$	$\neg B$

You now want, on line  $k$ , to prove  $\neg A$ . You can do this with the rule of MT, introduced in §27. But equally well, you can do this with the *basic* rules of §26:

$m$	$A \rightarrow B$	
$n$	$\neg B$	
$k$	$A$	
$k+1$	$B$	$\rightarrow E\ m, k$
$k+2$	$\perp$	$\neg E\ k+1, n$
$k+3$	$\neg A$	$\neg I\ k-k+2$

Again, the rule of MT can be derived from the *basic* rules of §26.

### 30.4 Derivation of De Morgan rules

Here is a derivation of the first De Morgan rule, from our basic rules:

$m$	$\neg(A \wedge B)$							
$k$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>A</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td></td> </tr> </table>	$A$						
$A$								
$k+1$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"></td> <td style="border-left: 1px solid black; padding-left: 5px;"><math>B</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td style="border-left: 1px solid black; border-bottom: 1px solid black; padding-left: 5px;"></td> <td></td> </tr> </table>		$B$					
	$B$							
$k+2$	$A \wedge B$	$\wedge I\ k, k+1$						
$k+3$	$\perp$	$\neg E\ k+2, m$						
$k+4$	$\neg B$	$\neg I\ k+1-k+3$						
$k+5$	$\neg A \vee \neg B$	$\vee I\ k+4$						
$k+6$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>\neg A</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td></td> </tr> </table>	$\neg A$						
$\neg A$								
$k+7$	$\neg A \vee \neg B$	$\vee I\ k+6$						
$k+8$	$\neg A \vee \neg B$	TND $k-k+5, k+6-k+7$						

Here is a derivation of the second De Morgan rule:

$m$	$\neg A \vee \neg B$					
$k$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>A \wedge B</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td></td> </tr> </table>	$A \wedge B$				
$A \wedge B$						
$k+1$	$A$	$\wedge E\ k$				
$k+2$	$B$	$\wedge E\ k$				
$k+3$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>\neg A</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td></td> </tr> </table>	$\neg A$				
$\neg A$						
$k+4$	$\perp$	$\neg E\ k+1, k+3$				
$k+5$	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding-right: 5px;"><math>\neg B</math></td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; border-bottom: 1px solid black; padding-right: 5px;"></td> <td></td> </tr> </table>	$\neg B$				
$\neg B$						
$k+6$	$\perp$	$\neg E\ k+2, k+5$				
$k+7$	$\perp$	$\vee E\ m, k+3-k+4, k+5-k+6$				
$k+8$	$\neg(A \wedge B)$	$\neg I\ k-k+7$				

Similar demonstrations can be offered explaining how we could derive the third and fourth De Morgan rules. But I will leave these as exercises.

### 30.5 Derivation of Double-negation elimination

Here is a derivation of the DNE rule, using only the *basic* rules.

$m$	$\neg\neg A$	
$k$	$A$	
$k+1$	$A$	R $k$
$k+2$	$\neg A$	
$k+3$	$\perp$	$\neg$ E $k+2, m$
$k+4$	$A$	X $k+3$
$k+5$	$A$	TND $k-k+1, k+2-k+4$

### Practice exercises

A. Provide proof schemes that justify the addition of the third and fourth De Morgan rules as derived rules.

B. The proofs you offered in response to the practice exercises of §§27–28 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.)

7

## Natural deduction for FOL

# Basic rules for FOL

The language of FOL uses of all of the connectives of TFL. So proofs in FOL 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

From the claim that everything is F, you can infer that any particular thing is F. You name it; it’s F. So the following should be fine:

1	$\forall xRxxd$	
2	$Raad$	$\forall E$ 1

We obtained line 2 by dropping the universal quantifier and replacing every instance of ‘ $x$ ’ with ‘ $a$ ’. Equally, the following should be allowed:

1	$\forall xRxxd$	
2	$Rddd$	$\forall E$ 1

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 ( $\forall E$ ):

$m$	$\forall xA(\dots x \dots x \dots)$	
	$A(\dots c \dots c \dots)$	$\forall E$ $m$

The notation here was introduced in §21. The point is that you can obtain any *substitution instance* of a universally quantified formula: replace every instance of the quantified variable with 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 logical operator. So the following is *banned*:

1	$\forall xBx \rightarrow Bk$	
2	$Bb \rightarrow Bk$	naughtily attempting to invoke $\forall E$ 1

This is illegitimate, since ‘ $\forall x$ ’ is not the main logical operator in line 1. (If you need a reminder as to why this sort of inference should be banned, reread §15.)

## 31.2 Existential introduction

From the claim that some particular thing is an F, you can infer that something is an F. So we ought to allow:

$$\begin{array}{l|l} 1 & Raad \\ \hline 2 & \exists xRaax \quad \exists I 1 \end{array}$$

Here, we have replaced the name ‘ $d$ ’ with a variable ‘ $x$ ’, and then existentially quantified over it. Equally, we would have allowed:

$$\begin{array}{l|l} 1 & Raad \\ \hline 2 & \exists xRxxd \quad \exists I 1 \end{array}$$

Here we have replaced both instances of the name ‘ $a$ ’ with a variable, and then existentially generalised. But we do not need to replace *both* instances of a name with a variable: if Narcissus loves himself, then there is someone who loves Narcissus. So we also allow:

$$\begin{array}{l|l} 1 & Raad \\ \hline 2 & \exists xRxad \quad \exists I 1 \end{array}$$

Here we have replaced *one* instance of the name ‘ $a$ ’ with a variable, and then existentially generalised. These observations motivate our introduction rule, although to explain it, we shall need to introduce some new notation.

Where  $A$  is a sentence containing the name  $c$ , we can emphasise this by writing ‘ $A(\dots c \dots c \dots)$ ’. We shall write ‘ $A(\dots x \dots c \dots)$ ’ to indicate any formula obtained by replacing *some or all* of the instances of the name  $c$  with the variable  $x$ . Armed with this, our introduction rule is:

$\begin{array}{l l} m & A(\dots c \dots c \dots) \\ \hline & \exists xA(\dots x \dots c \dots) \quad \exists I m \end{array}$
$x$ must not occur in $A(\dots c \dots c \dots)$

The constraint is included to guarantee that any application of the rule yields a sentence of FOL. Thus the following is allowed:

$$\begin{array}{l|l} 1 & Raad \\ \hline 2 & \exists xRxad \quad \exists I 1 \\ 3 & \exists y\exists xRxyd \quad \exists I 2 \end{array}$$

But this is banned:

1	$Raad$	
2	$\exists xRxad$	$\exists I$ 1
3	$\exists x\exists xRxxd$	naughtily attempting to invoke $\exists I$ 2

since the expression on line 3 contains clashing variables, and so is not a sentence of FOL.

### 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 §14, we stipulated that domains in FOL must have at least one member. We then defined a logical truth (of FOL) 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, 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 FOL, 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 FOL, 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 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:

$$\begin{array}{l|l} 1 & \forall xFx \\ \hline 2 & Fa \quad \forall E 1 \end{array}$$

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 there 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:

$$\begin{array}{l|l} 1 & \forall xFx \\ \hline 2 & Fa \quad \forall E 1 \\ 3 & \forall yFy \quad \forall I 2 \end{array}$$

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 ( $\forall I$ ):

$\begin{array}{l l} m & A(\dots c \dots c \dots) \\ \hline & \forall x A(\dots x \dots x \dots) \quad \forall I m \end{array}$ <p><math>c</math> must not occur in any undischarged assumption</p> <p><math>x</math> must not occur in <math>A(\dots c \dots c \dots)</math></p>
--

A crucial aspect of this rule, though, is bound up in the first constraint. This constraint ensures that we are always reasoning at a sufficiently general level.<sup>1</sup> To see the constraint in action, consider this terrible argument:

<sup>1</sup>Recall from §26 that we are treating ' $\perp$ ' as a canonical contradiction. But if it were the canonical contradiction as involving some *constant*, it might interfere with the constraint mentioned here. To avoid such problems, we shall treat ' $\perp$ ' as a canonical contradiction *that involves no particular names*.

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, then we are not reasoning generally enough to license  $\forall I$ .

Although the name may not occur in any *undischarged* assumption, it may occur in a *discharged* assumption. That is, it may occur in a subproof that we have already closed. For example, this is just fine:

1	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>Gd</math></td> </tr> </table>	$Gd$	
$Gd$			
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.

I should emphasise one last point. As per the conventions of §21.3, the use of  $\forall I$  requires that we are replacing *every* instance of the name  $c$  in  $A(\dots x \dots x \dots)$  with the variable  $x$ . If we only replace *some* names and not others, we end up ‘proving’ silly things. For example, consider the argument:

Everyone is as old as themselves; so everyone is as old as Judi Dench

We might symbolise this as follows:

$$\forall xOxx \therefore \forall xOxd$$

But now suppose we tried to *vindicate* this terrible argument with the following:

1	$\forall xOxx$	
2	$Odd$	$\forall E$ 1
3	$\forall xOxd$	naughtily attempting to invoke $\forall I$ 2

Fortunately, our rules do not allow for us to do this: the attempted proof is banned, since it doesn’t replace *every* occurrence of ‘ $d$ ’ in line 2 with an ‘ $x$ ’.

### 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 ( $\exists E$ ):

$m$	$\exists xA(\dots x \dots x \dots)$	
$i$	$A(\dots c \dots c \dots)$	
$j$	$B$	
	$B$	$\exists E$ $m, i-j$
$c$ must not occur in any assumption undischarged before line $i$ $c$ must not occur in $\exists xA(\dots x \dots x \dots)$ $c$ 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. Someone 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	$\exists x \neg Lx$					
3	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>\neg Lb</math></td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>Lb \wedge \neg Lb</math></td> <td style="padding-left: 10px;"><math>\wedge E</math> 1, 3</td> </tr> </table>	$\neg Lb$		$Lb \wedge \neg Lb$	$\wedge E$ 1, 3	
$\neg Lb$						
$Lb \wedge \neg Lb$	$\wedge E$ 1, 3					
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 this would be no better:

1	Lb							
2	$\exists x \neg Lx$							
3	<table style="border-collapse: collapse; margin-left: 1em;"> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>\neg Lb</math></td> <td></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>Lb \wedge \neg Lb</math></td> <td style="padding-left: 10px;"><math>\wedge E</math> 1, 3</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 5px;"><math>\exists x(Lx \wedge \neg Lx)</math></td> <td style="padding-left: 10px;"><math>\exists I</math> 4</td> </tr> </table>	$\neg Lb$		$Lb \wedge \neg Lb$	$\wedge E$ 1, 3	$\exists x(Lx \wedge \neg Lx)$	$\exists I$ 4	
$\neg Lb$								
$Lb \wedge \neg Lb$	$\wedge E$ 1, 3							
$\exists x(Lx \wedge \neg Lx)$	$\exists I$ 4							
6	$\exists x(Lx \wedge \neg Lx)$	naughtily attempting to invoke $\exists E$ 2, 3–5						

The last line is 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 existential quantifier, choose a new name for your substitution instance.* That way, you can guarantee that you meet all the constraints on the rule for  $\exists E$ .

## Practice exercises

A. Explain why these two 'proofs' are incorrect. Also, provide interpretations invalidate the fallacious argument forms the 'proofs' enshrine:

1	$\forall x Rxx$		1	$\forall x \exists y Rxy$					
2	Raa	$\forall E$ 1	2	$\exists y Ray$	$\forall E$ 1				
3	$\forall y Ray$	$\forall I$ 2	3	<table style="border-collapse: collapse; margin-left: 1em;"> <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;"><math>\exists x Rxx</math></td> <td style="padding-left: 10px;"><math>\exists I</math> 3</td> </tr> </table>	Raa		$\exists x Rxx$	$\exists I$ 3	
Raa									
$\exists x Rxx$	$\exists I$ 3								
4	$\forall x \forall y Rxy$	$\forall I$ 3	5	$\exists x Rxx$	$\exists E$ 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	$\forall x \neg Rmx$				
3	$\exists y (Rmy \vee Rym)$				
4	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Rma \vee Ram</math></td> </tr> <tr> <td style="border-top: 1px solid black; border-left: 1px solid black; padding-left: 10px;"><math>\neg Rma</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Ram</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>\exists x Rxm</math></td> </tr> </table>	$Rma \vee Ram$	$\neg Rma$	$Ram$	$\exists x Rxm$
$Rma \vee Ram$					
$\neg Rma$					
$Ram$					
$\exists x Rxm$					
8	$\exists x Rxm$				

1	$\forall x (\exists y Lxy \rightarrow \forall z Lzx)$
2	$Lab$
3	$\exists y Lay \rightarrow \forall z Lza$
4	$\exists y Lay$
5	$\forall z Lza$
6	$Lca$
7	$\exists y Lcy \rightarrow \forall z Lzc$
8	$\exists y Lcy$
9	$\forall z Lzc$
10	$Lcc$
11	$\forall x Lxx$

1	$\forall x (\exists x \rightarrow Kx)$							
2	$\exists x \forall y Lxy$							
3	$\forall x \exists x$							
4	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>\forall y Lay</math></td> </tr> <tr> <td style="border-top: 1px solid black; border-left: 1px solid black; padding-left: 10px;"><math>Laa</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Ja</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Ja \rightarrow Ka</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Ka</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>Ka \wedge Laa</math></td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;"><math>\exists x (Kx \wedge Lxx)</math></td> </tr> </table>	$\forall y Lay$	$Laa$	$Ja$	$Ja \rightarrow Ka$	$Ka$	$Ka \wedge Laa$	$\exists x (Kx \wedge Lxx)$
$\forall y Lay$								
$Laa$								
$Ja$								
$Ja \rightarrow Ka$								
$Ka$								
$Ka \wedge Laa$								
$\exists x (Kx \wedge Lxx)$								
11	$\exists x (Kx \wedge Lxx)$							

C. In §15 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 these argument forms in FOL 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 FOL, 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.  $\forall xPx \rightarrow \forall xQx, \exists x\neg Px \therefore \exists x\neg Qx$

# Conversion of quantifiers

32

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.

In §14, we noted that  $\neg\exists xA$  is logically equivalent to  $\forall x\neg A$ . We shall add some rules to our proof system that govern this. In particular, we add:

$m$		$\forall x\neg A$	
		$\neg\exists xA$	CQ $m$

and

$m$		$\neg\exists xA$	
		$\forall x\neg A$	CQ $m$

Equally, we add:

$m$		$\exists x\neg A$	
		$\neg\forall xA$	CQ $m$

and

$m$		$\neg\forall xA$	
		$\exists x\neg A$	CQ $m$

## Practice exercises

A. Show that the following are jointly contrary:

1.  $Sa \rightarrow Tm, Tm \rightarrow Sa, Tm \wedge \neg Sa$
2.  $\neg\exists xRxa, \forall x\forall yRyx$
3.  $\neg\exists x\exists yLxy, Laa$
4.  $\forall x(Px \rightarrow Qx), \forall z(Pz \rightarrow Rz), \forall yPy, \neg Qa \wedge \neg Rb$

B. Show that each pair of sentences is provably equivalent:

1.  $\forall x(Ax \rightarrow \neg Bx), \neg \exists x(Ax \wedge Bx)$
2.  $\forall x(\neg Ax \rightarrow Bd), \forall xAx \vee Bd$

C. In §15, I considered what happens when we move quantifiers ‘across’ various logical operators. Show that each pair of sentences is provably equivalent:

1.  $\forall x(Fx \wedge Ga), \forall xFx \wedge Ga$
2.  $\exists x(Fx \vee Ga), \exists xFx \vee Ga$
3.  $\forall x(Ga \rightarrow Fx), Ga \rightarrow \forall xFx$
4.  $\forall x(Fx \rightarrow Ga), \exists xFx \rightarrow Ga$
5.  $\exists x(Ga \rightarrow Fx), Ga \rightarrow \exists xFx$
6.  $\exists x(Fx \rightarrow Ga), \forall xFx \rightarrow Ga$

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.

# Rules for identity

In §20, I mentioned the philosophically contentious thesis of the *identity of indiscernibles*. This is the claim that objects which are indiscernible in every way are, in fact, identical to each other. I also mentioned that we will not subscribe to this thesis. It follows that, no matter how much you tell me about two objects, I cannot prove that they are identical. Unless, of course, you tell me that the two objects are, in fact, identical. But then the proof will hardly be very illuminating.

The general point, though, is that *no sentences* which do not already contain the identity predicate could justify an inference to ' $a = b$ '. So our identity introduction rule cannot allow us to infer to an 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:

$\begin{array}{l}   \\ \hline \mathbf{c} = \mathbf{c} \quad =I \end{array}$
---

Notice that this rule does not require referring to any prior lines of the proof. For any name  $\mathbf{c}$ , you can write  $\mathbf{c} = \mathbf{c}$  on 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 ' $b$ '. For any sentence with ' $a$ ' in it, you can replace some or all of the occurrences of ' $a$ ' with ' $b$ ' and produce an equivalent sentence. For example, from ' $Raa$ ' and ' $a = b$ ', you are justified in inferring ' $Rab$ ', ' $Rba$ ' or ' $Rbb$ '. More generally:

$\begin{array}{l} m \\ n \end{array} \begin{array}{l}   \\ \hline \mathbf{a} = \mathbf{b} \\ \mathbf{A}(\dots\mathbf{a}\dots\mathbf{a}\dots) \\ \mathbf{A}(\dots\mathbf{b}\dots\mathbf{a}\dots) \end{array} \quad =E \ m, n$
--

The notation here is as for  $\exists$ I. So  $\mathbf{A}(\dots\mathbf{a}\dots\mathbf{a}\dots)$  is a formula containing the name  $\mathbf{a}$ , and  $\mathbf{A}(\dots\mathbf{b}\dots\mathbf{a}\dots)$  is a formula obtained by replacing one or more instances of the name  $\mathbf{a}$  with the name  $\mathbf{b}$ . 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. Symmetrically, we allow:

$m$	$a = b$	
$n$	$A(\dots b \dots b \dots)$	
	$A(\dots a \dots b \dots)$	$=E\ m, n$

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	$\forall x \forall y(x = y \rightarrow y = x)$	$\forall I\ 5$

We obtain line 3 by replacing one instance of ' $a$ ' in line 2 with an instance of ' $b$ '; this is justified given ' $a = b$ '.

Second, we shall prove that identity is *transitive*:

1	$a = b \wedge b = c$	
2	$a = b$	$\wedge E\ 1$
3	$b = c$	$\wedge E\ 1$
4	$a = c$	$=E\ 2, 3$
5	$(a = b \wedge b = c) \rightarrow a = c$	$\rightarrow I\ 1-4$
6	$\forall z((a = b \wedge b = z) \rightarrow a = z)$	$\forall I\ 5$
7	$\forall y \forall z((a = y \wedge y = z) \rightarrow a = z)$	$\forall I\ 6$
8	$\forall x \forall y \forall z((x = y \wedge y = z) \rightarrow x = z)$	$\forall I\ 7$

We obtain line 4 by replacing ' $b$ ' in line 3 with ' $a$ '; this is justified given ' $a = b$ '.

## 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.  $\exists x \neg \neg x, \exists x \neg \neg x \vdash \exists x \exists y \neg x = y$
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 §17, 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.

# Derived rules

As in the case of TFL, I first introduced some rules for FOL as basic (in §31), and then added some further rules for conversion of quantifiers (in §32). In fact, the CQ rules should be regarded as *derived* rules, for they can be derived from the *basic* rules of §31. (The point here is as in §30.) Here is a justification for the first CQ rule:

$m$	$\forall x \neg Ax$	
$k$	$\exists x Ax$	
$k+1$	$A c$	
$k+2$	$\neg Ac$	$\forall E\ m$
$k+3$	$\perp$	$\neg E\ k+1, k+2$
$k+4$	$\perp$	$\exists E\ k, k+1-k+3$
$k+5$	$\neg \exists x Ax$	$\neg I\ k-k+4$

Here is a justification of the second CQ rule:

$m$	$\exists x \neg Ax$	
$k$	$\forall x Ax$	
$k+1$	$\neg Ac$	
$k+2$	$Ac$	$\forall E\ k$
$k+3$	$\perp$	$\neg E\ k+2, k+1$
$k+4$	$\perp$	$\exists E\ m, k+1-k+3$
$k+5$	$\neg \forall x Ax$	$\neg I\ k-k+4$

This explains why the CQ rules can be treated as derived. Similar justifications can be offered for the other two CQ rules.

## Practice exercises

A. Offer proofs which justify the addition of the third and fourth CQ rules as derived rules.

# Proof-theoretic concepts and semantic concepts

35

We have used two different turnstiles in this book. This:

$$A_1, A_2, \dots, A_n \vdash C$$

means that there is some proof which ends with  $C$  and whose only undischarged assumptions are among  $A_1, A_2, \dots, A_n$ . This is a *proof-theoretic notion*. By contrast, this:

$$A_1, A_2, \dots, A_n \models C$$

means that no valuation (or interpretation) makes all of  $A_1, A_2, \dots, A_n$  true and  $C$  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*.

At the risk of repetition: *They are different notions*.

Once you have fully 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  $A$  on no assumptions, and thus show that  $A$  is a theorem, i.e.  $\vdash A$ , we can legitimately infer that  $A$  is a logical truth, i.e.,  $\models A$ . Similarly, if we construct a model in which  $A$  is false and thus show that it is not a logical truth, i.e.  $\not\models A$ , it follows that  $A$  is not a theorem, i.e.  $\not\vdash A$ .

More generally, we have the following powerful result:

$$A_1, A_2, \dots, A_n \vdash B \text{ iff } A_1, A_2, \dots, A_n \models B$$

This shows that, whilst provability and entailment are *different* notions, they are extensionally equivalent. 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.

In fact, demonstrating that provability and semantic entailment agree is, very decisively, the point at which introductory logic becomes intermediary logic. Agreement, in the case of TFL, is covered in a little sequel to this book, *Metatheory*. Agreement, in the case of FOL, is one of the first big results in mathematical logic.

	Yes	No
Is $A$ a logical truth?	give a proof which shows $\vdash A$	give an interpretation in which $A$ is false
Is $A$ a contradiction?	give a proof which shows $\vdash \neg A$	give an interpretation in which $A$ is true
Are $A$ and $B$ equivalent?	give two proofs, one for $A \vdash B$ and one for $B \vdash A$	give an interpretation in which $A$ and $B$ have different truth values
Are $A_1, A_2, \dots, A_n$ jointly consistent?	give an interpretation in which all of $A_1, A_2, \dots, A_n$ are true	prove a contradiction from assumptions $A_1, A_2, \dots, A_n$
Is $A_1, A_2, \dots, A_n \therefore C$ valid?	give a proof with assumptions $A_1, A_2, \dots, A_n$ and concluding with $C$	give an interpretation in which each of $A_1, A_2, \dots, A_n$ is true and $C$ is false

# Appendices

## 36.1 Alternative nomenclature

**Truth-functional logic.** TFL goes by other names. Sometimes it is called *sentence logic*, because it deals fundamentally with sentences. Sometimes it is called *propositional logic*, on the idea that it deals fundamentally with propositions. I have stuck with *truth-functional logic*, to emphasise the fact that it deals only with assignments of truth and falsity to sentences, and that its connectives are all truth-functional.

**First-order logic.** FOL goes by other names. Sometimes it is called *predicate logic*, because it allows us to apply predicates to objects. Sometimes it is called *quantified logic*, because it makes use of quantifiers.

**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 countenance ‘ill-formed formulas’). I have stuck with ‘formula’.

In §6, I defined *sentences* of TFL. These are also sometimes called ‘formulas’ (or ‘well-formed formulas’) since in TFL, unlike FOL, there is no distinction between a formula and a sentence.

**Valuations.** Some texts call valuations *truth-assignments*.

**Expressive adequacy.** Some texts describe TFL 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’, ‘trinary’, etc.

**Names.** In FOL, I have used ‘*a*’, ‘*b*’, ‘*c*’, for names. Some texts call these ‘constants’. 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’.

## 36.2 Alternative symbols

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 *hoe*, ‘ $\neg$ ’, and the *swung dash*, ‘ $\sim$ ’. In some more advanced formal systems it is necessary to distinguish between two kinds of negation; the distinction is sometimes represented by using both ‘ $\neg$ ’ and ‘ $\sim$ ’. Some texts use ‘ $x \neq y$ ’ to abbreviate ‘ $\neg x = y$ ’.

**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.) 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’.

**Material Conditional.** There are two common symbols for the material conditional: the *arrow*, ‘ $\rightarrow$ ’, and the *hook*, ‘ $\supset$ ’.

**Material Biconditional.** The *double-headed arrow*, ‘ $\leftrightarrow$ ’, is used in systems that use the arrow to represent the material conditional. Systems that use the hook for the conditional typically use the *triple bar*, ‘ $\equiv$ ’, for the biconditional.

**Quantifiers.** The universal quantifier is typically 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 typographies are summarised below:

negation	$\neg, \sim$
conjunction	$\wedge, \&, \cdot$
disjunction	$\vee$
conditional	$\rightarrow, \supset$
biconditional	$\leftrightarrow, \equiv$
universal quantifier	$\forall x, (x)$

In formulating my natural deduction system, I treated certain rules of natural deduction as *basic*, and others as *derived*. However, I could equally well have taken various different rules as basic or derived. I shall illustrate this point by considering some alternative treatments of disjunction, negation, and the quantifiers. I shall also explain why I have made the choices that I have.

## 37.1 Alternative negation rules

Some systems take the following rule as their basic negation introduction rule:

$$\begin{array}{l}
 m \\
 n - 1 \\
 n
 \end{array}
 \left| \begin{array}{l}
 \hline
 A \\
 B \\
 \hline
 \neg B
 \end{array} \right.
 \begin{array}{l}
 \\
 \\
 \neg A
 \end{array}
 \quad \neg I^* \ m-n$$

and the following as their basic negation elimination rule:

$$\begin{array}{l}
 m \\
 n - 1 \\
 n
 \end{array}
 \left| \begin{array}{l}
 \neg A \\
 \hline
 B \\
 \neg B
 \end{array} \right.
 \begin{array}{l}
 \\
 \\
 A
 \end{array}
 \quad \neg E^* \ m-n$$

Using these two rules, we could have derived all of the rules governing negation and contradiction that we have taken as basic (i.e.  $\neg I$ ,  $\neg E$ , X and TND). Indeed, we could have avoided all use of the symbol ' $\perp$ ' altogether. Negation would have had a single introduction and elimination rule, and would have behaved much more like the other connectives.

The resulting system would have had fewer rules than ours. So why did I chose to separate out contradiction, and to use an explicit rule TND?<sup>1</sup>

My first reason is that adding the symbol ' $\perp$ ' to our natural deduction system makes proofs considerably easier to work with.

My second reason is that a lot of fascinating philosophical discussion has focussed on the acceptability or otherwise of *ex falso quodlibet* (i.e. X) and *tertium non datur* (i.e. TND). By treating these as separate rules in the proof system, we will be in a better position

<sup>1</sup>P.D. Magnus's original version of this book went the other way.

to engage with that philosophical discussion. In particular: having invoked these rules explicitly, it will be much easier for us to know what a system which lacked these rules would look like.

### 37.2 Alternative disjunction elimination

Some systems take DS as their basic rule for disjunction elimination. They can then treat the  $\vee E$  rule as a derived rule, by using the following proof scheme:

$m$	$A \vee B$	
$i$	$A$	
$j$	$C$	
$k$	$B$	
$l$	$C$	
$n$	$A \rightarrow C$	$\rightarrow I\ i-j$
$n+1$	$B \rightarrow C$	$\rightarrow I\ k-l$
$n+2$	$C$	
$n+3$	$C$	$R\ n+2$
$n+4$	$\neg C$	
$n+5$	$A$	
$n+6$	$C$	$\rightarrow E\ n, n+5$
$n+7$	$\perp$	$\neg E\ n+6, n+4$
$n+8$	$\neg A$	$\neg I\ n+5-n+7$
$n+9$	$B$	$DS\ m, n+8$
$n+10$	$C$	$\rightarrow E\ n+1, n+9$
$n+11$	$C$	$TND\ n+2-n+3, n+4-n+10$

So why did I choose to take  $\vee E$  as basic, rather than DS?<sup>2</sup> My reasoning is that DS involves the use of ‘ $\neg$ ’ in the statement of the rule. It is in some sense ‘cleaner’ for our disjunction elimination rule to avoid mentioning *other* connectives.

### 37.3 Alternative quantification rules

An alternative approach to the quantifiers is to take as basic the rules for  $\forall I$  and  $\forall E$  from §31, and also two CQ rule which allow us to move from  $\forall x\neg A$  to  $\neg\exists xA$  and vice versa.<sup>3</sup>

<sup>2</sup>P.D. Magnus’s original version of this book went the other way.

<sup>3</sup>Warren Goldfarb follows this line in *Deductive Logic*, 2003, Hackett Publishing Co.

Taking only these rules as basic, we could have derived the  $\exists$ I and  $\exists$ E rules provided in §31. To derive the  $\exists$ I rule is fairly simple. Suppose  $A$  contains the name  $c$ , and contains no instances of the variable  $x$ , and that we want to do the following:

$$\begin{array}{l|l} m & A(\dots c \dots c \dots) \\ k & \exists x A(\dots x \dots c \dots) \end{array}$$

This is not yet permitted, since in this new system, we do not have the  $\exists$ I rule. We can, however, offer the following:

$$\begin{array}{l|l|l} m & A(\dots c \dots c \dots) & \\ m+1 & \frac{}{\neg \exists x A(\dots x \dots c \dots)} & \\ m+2 & \frac{}{\forall x \neg A(\dots x \dots c \dots)} & \text{CQ } m+1 \\ m+3 & \neg A(\dots c \dots c \dots) & \forall E \ m+2 \\ m+4 & \perp & \neg E \ m, m+3 \\ m+5 & \neg \neg \exists x A(\dots x \dots c \dots) & \neg I \ m+1-m+4 \\ m+6 & \exists x A(\dots x \dots c \dots) & DNE \ m+5 \end{array}$$

To derive the  $\exists$ E rule is rather more subtle. This is because the  $\exists$ E rule has an important constraint (as, indeed, does the  $\forall$ I rule), and we need to make sure that we are respecting it. So, suppose we are in a situation where we *want* to do the following:

$$\begin{array}{l|l|l} m & \exists x A(\dots x \dots x \dots) & \\ i & \frac{}{A(\dots c \dots c \dots)} & \\ j & \frac{}{B} & \\ k & B & \end{array}$$

where  $c$  does not occur in any undischarged assumptions, or in  $B$ , or in  $\exists x A(\dots x \dots x \dots)$ . Ordinarily, we would be allowed to use the  $\exists$ E rule; but we are not here assuming that we have access to this rule as a basic rule. Nevertheless, we could offer the following, more complicated derivation:

$m$	$\exists \mathbf{x}A(\dots \mathbf{x} \dots \mathbf{x} \dots)$	
$i$	$A(\dots \mathbf{c} \dots \mathbf{c} \dots)$	
$j$	$\mathbf{B}$	
$k$	$A(\dots \mathbf{c} \dots \mathbf{c} \dots) \rightarrow \mathbf{B}$	$\rightarrow I\ i-j$
$k+1$	$\neg \mathbf{B}$	
$k+2$	$\neg A(\dots \mathbf{c} \dots \mathbf{c} \dots)$	$MT\ k, k+1$
$k+3$	$\forall \mathbf{x} \neg A(\dots \mathbf{x} \dots \mathbf{x} \dots)$	$\forall I\ k+2$
$k+4$	$\neg \exists \mathbf{x} A(\dots \mathbf{x} \dots \mathbf{x} \dots)$	$CQ\ k+3$
$k+5$	$\perp$	$\neg E\ m, k+4$
$k+6$	$\neg \neg \mathbf{B}$	$\neg I\ k+1-k+5$
$k+7$	$\mathbf{B}$	$DNE\ k+6$

We are permitted to use  $\forall I$  on line  $k+3$  because  $\mathbf{c}$  does not occur in any undischarged assumptions or in  $\mathbf{B}$ . The entries on lines  $k+4$  and  $k+1$  contradict each other, because  $\mathbf{c}$  does not occur in  $\exists \mathbf{x}A(\dots \mathbf{x} \dots \mathbf{x} \dots)$ .

Armed with these derived rules, we could now go on to derive the two remaining CQ rules, exactly as in §34.

So, why did I start with all of the quantifier rules as basic, and then derive the CQ rules?

My first reason is that it seems more intuitive to treat the quantifiers as on a par with one another, giving them their own basic rules for introduction and elimination.

My second reason relates to the discussion of alternative negation rules. In the derivations of the rules of  $\exists I$  and  $\exists E$  that I have offered in this section, I invoked DNE. This is a derived rule, whose derivation essentially depends upon the use of TND. But, as I mentioned earlier, TND is a contentious rule. So, if we want to move to a system which abandons TND, but which still allows us to use existential quantifiers, we shall want to take the introduction and elimination rules for the quantifiers as basic, and take the CQ rules as derived. (Indeed, in a system without TND, we shall be *unable* to derive the CQ rule which moves from  $\neg \forall \mathbf{x}A$  to  $\exists \mathbf{x} \neg A$ .)

# Quick reference

## 38.1 Characteristic Truth Tables

$A$	$\neg A$	$A$	$B$	$A \wedge B$	$A \vee B$	$A \rightarrow B$	$A \leftrightarrow B$
T	F	T	T	T	T	T	T
F	T	T	F	F	T	F	F
		F	T	F	T	T	F
		F	F	F	F	T	T

## 38.2 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	$\exists x(Fx \wedge \forall y(Fy \rightarrow x = y) \wedge \neg Gx)$

### 38.3 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:  $\exists x_1\exists x_2\exists x_3(Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3)$

four:  $\exists x_1\exists x_2\exists x_3\exists x_4(Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge Fx_4 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_1 = x_4 \wedge \neg x_2 = x_3 \wedge \neg x_2 = x_4 \wedge \neg x_3 = x_4)$

$n$ :  $\exists x_1\dots\exists x_n(Fx_1 \wedge \dots \wedge Fx_n \wedge \neg x_1 = x_2 \wedge \dots \wedge \neg x_{n-1} = x_n)$

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:  $\exists x_1\exists x_2\exists x_3[Fx_1 \wedge Fx_2 \wedge Fx_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3 \wedge \forall y(Fy \rightarrow (y = x_1 \vee y = x_2 \vee y = x_3))]$

$n$ :  $\exists x_1\dots\exists x_n[Fx_1 \wedge \dots \wedge Fx_n \wedge \neg x_1 = x_2 \wedge \dots \wedge \neg x_{n-1} = x_n \wedge \forall y(Fy \rightarrow (y = x_1 \vee \dots \vee y = x_n))]$

## 38.4 Basic deduction rules for TFL

## Conjunction

$$\begin{array}{l|l}
 m & A \\
 n & B \\
 & \hline
 & A \wedge B \quad \wedge I \ m, n
 \end{array}$$

$$\begin{array}{l|l}
 m & A \wedge B \\
 & \hline
 & A \quad \wedge E \ m
 \end{array}$$

$$\begin{array}{l|l}
 m & A \wedge B \\
 & \hline
 m & B \quad \wedge E \ m
 \end{array}$$

## Disjunction

$$\begin{array}{l|l}
 m & A \\
 & \hline
 & A \vee B \quad \vee I \ m
 \end{array}$$

$$\begin{array}{l|l}
 m & A \\
 & \hline
 m & B \vee A \quad \vee I \ m
 \end{array}$$

$$\begin{array}{l|l}
 m & A \vee B \\
 i & \begin{array}{l|l} A \\ \hline C \end{array} \\
 j & \begin{array}{l|l} B \\ \hline C \end{array} \\
 k & \begin{array}{l|l} B \\ \hline C \end{array} \\
 l & \begin{array}{l|l} C \\ \hline C \end{array} \\
 & \hline
 & C \quad \vee E \ m, i-j, k-l
 \end{array}$$

## Conditional

$$\begin{array}{l|l}
 i & \begin{array}{l|l} A \\ \hline B \end{array} \\
 j & \begin{array}{l|l} B \\ \hline A \rightarrow B \end{array} \\
 & \hline
 & A \rightarrow B \quad \rightarrow I \ i-j
 \end{array}$$

$$\begin{array}{l|l}
 m & A \rightarrow B \\
 n & A \\
 & \hline
 n & B \quad \rightarrow E \ m, n
 \end{array}$$

## Biconditional

$$\begin{array}{l|l}
 i & \begin{array}{l|l} A \\ \hline B \end{array} \\
 j & \begin{array}{l|l} B \\ \hline A \end{array} \\
 k & \begin{array}{l|l} B \\ \hline A \end{array} \\
 l & \begin{array}{l|l} A \\ \hline A \leftrightarrow B \end{array} \\
 & \hline
 & A \leftrightarrow B \quad \leftrightarrow I \ i-j, k-l
 \end{array}$$

$$\begin{array}{l|l}
 m & A \leftrightarrow B \\
 n & A \\
 & \hline
 n & B \quad \leftrightarrow E \ m, n
 \end{array}$$

$$\begin{array}{l|l}
 m & A \leftrightarrow B \\
 n & B \\
 & \hline
 n & A \quad \leftrightarrow E \ m, n
 \end{array}$$

## Negation and contradiction

$$\begin{array}{l|l}
 m & A \\
 n & \neg A \\
 & \hline
 & \perp \quad \neg E \ m, n
 \end{array}$$

$$\begin{array}{l|l}
 i & \begin{array}{l|l} A \\ \hline \perp \end{array} \\
 j & \begin{array}{l|l} \perp \\ \hline \neg A \end{array} \\
 & \hline
 & \neg A \quad \neg I \ i-j
 \end{array}$$

$$\begin{array}{l|l}
 m & \perp \\
 & \hline
 m & A \quad X \ m \\
 i & \begin{array}{l|l} A \\ \hline B \end{array} \\
 j & \begin{array}{l|l} B \\ \hline \neg A \end{array} \\
 k & \begin{array}{l|l} \neg A \\ \hline B \end{array} \\
 l & \begin{array}{l|l} B \\ \hline B \end{array} \\
 & \hline
 & B \quad \text{TND } i-j, k-l
 \end{array}$$

### 38.5 Derived rules for TFL

#### Reiteration

$$\begin{array}{l|l}
 m & A \\
 & A \quad R\ m
 \end{array}$$

#### Double-negation elimination

$$\begin{array}{l|l}
 m & \neg\neg A \\
 & A \quad DNE\ m
 \end{array}$$

#### Disjunctive syllogism

$$\begin{array}{l|l}
 m & A \vee B \\
 n & \neg A \\
 & B \quad DS\ m, n
 \end{array}$$

#### De Morgan Rules

$$\begin{array}{l|l}
 m & \neg(A \vee B) \\
 & \neg A \wedge \neg B \quad DeM\ m
 \end{array}$$

$$\begin{array}{l|l}
 m & A \vee B \\
 n & \neg B \\
 & A \quad DS\ m, n
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg A \wedge \neg B \\
 & \neg(A \vee B) \quad DeM\ m
 \end{array}$$

#### Modus Tollens

$$\begin{array}{l|l}
 m & A \rightarrow B \\
 n & \neg B \\
 & \neg A \quad MT\ m, n
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg(A \wedge B) \\
 & \neg A \vee \neg B \quad DeM\ m
 \end{array}$$

$$\begin{array}{l|l}
 m & \neg A \vee \neg B \\
 & \neg(A \wedge B) \quad DeM\ m
 \end{array}$$

## 38.6 Basic deduction rules for FOL

### Universal elimination

$$m \left| \begin{array}{l} \forall x A(\dots x \dots x \dots) \\ A(\dots c \dots c \dots) \end{array} \right. \quad \forall E \ m$$

### Existential introduction

$$m \left| \begin{array}{l} A(\dots c \dots c \dots) \\ \exists x A(\dots x \dots c \dots) \end{array} \right. \quad \exists I \ m$$

$x$  must not occur in  $A(\dots c \dots c \dots)$

### Universal introduction

$$m \left| \begin{array}{l} A(\dots c \dots c \dots) \\ \forall x A(\dots x \dots x \dots) \end{array} \right. \quad \forall I \ m$$

$c$  must not occur in any undischarged assumption

$x$  must not occur in  $A(\dots c \dots c \dots)$

### Existential elimination

$$m \left| \begin{array}{l} \exists x A(\dots x \dots x \dots) \\ i \left| \begin{array}{l} A(\dots c \dots c \dots) \\ j \left| \begin{array}{l} B \end{array} \right. \\ B \end{array} \right. \end{array} \right. \quad \exists E \ m, i-j$$

$c$  must not occur in any undischarged assumption, in  $\exists x A(\dots x \dots x \dots)$ , or in  $B$

### Identity introduction

$$\left| \begin{array}{l} c = c \end{array} \right. \quad =I$$

### Identity elimination

$$m \left| \begin{array}{l} a = b \\ n \left| \begin{array}{l} A(\dots a \dots a \dots) \\ A(\dots b \dots a \dots) \end{array} \right. \end{array} \right. \quad =E \ m, n$$

$$m \left| \begin{array}{l} a = b \\ n \left| \begin{array}{l} A(\dots b \dots b \dots) \\ A(\dots a \dots b \dots) \end{array} \right. \end{array} \right. \quad =E \ m, n$$

## 38.7 Derived rules for FOL

$$m \left| \begin{array}{l} \forall x \neg A \\ \neg \exists x A \end{array} \right. \quad CQ \ m$$

$$m \left| \begin{array}{l} \exists x \neg A \\ \neg \forall x A \end{array} \right. \quad CQ \ m$$

$$m \left| \begin{array}{l} \neg \exists x A \\ \forall x \neg A \end{array} \right. \quad CQ \ m$$

$$m \left| \begin{array}{l} \neg \forall x A \\ \exists x \neg A \end{array} \right. \quad CQ \ m$$

In the Introduction to his volume *Symbolic Logic*, Charles Lutwidge Dodson advised: “When you come to any passage you don’t understand, *read it again*: if you *still* don’t understand it, *read it again*: if you fail, even after *three* readings, very likely your brain is getting a little tired. In that case, put the book away, and take to other occupations, and next day, when you come to it fresh, you will very likely find that it is *quite* easy.”

The same might be said for this volume, although readers are forgiven if they take a break for snacks after *two* readings.

P.D. Magnus is an associate professor of philosophy in Albany, New York. His primary research is in the philosophy of science.

Tim Button is a Senior Lecturer, and Fellow of St John’s College, at the University of Cambridge. He mostly works on meta(meta)physics, logic, mathematics, and language. You can find out more at [nottub.com](http://nottub.com).