forallχ

Saint Louis University 2017-18

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

Modified by:
Kathryn Lindeman
Saint Louis University

Typesetting was carried out entirely in LATEX2ε. The style for typesetting proofs is based on fitch.sty (v0.4) by Peter Selinger, University of Ottawa.

Contents

5   interpretations   96

6   natural deduction for TFL   117

7   Natural deduction for FOL   151

Appendices   172

Chapter 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 disagreements.

An argument, as we will 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 will 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 really pouring outside.
> And if you don't take an umbrella, you'll get soaked.

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

> It is really pouring outside, so you should take an umbrella. Since
> if you don't take an umbrella, you"ll get soaked.

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

> so, therefore, hence, thus, accordingly, consequently

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

> since, because, given that

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

Practice exercises

At the end of some sections, there are problems that review and explore the material covered in the chapter. There is no substitute for actually working through some problems, because logic is more about a way of thinking than it is about memorizing facts.

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

1. It is sunny. So I should take my sunglasses.
2. It must have been sunny. I did wear my sunglasses, after all.
3. No one but you has had their hands in the cookie-jar. And the scene of the crime is littered with cookie-crumbs. You're the culprit!
4. Miss Scarlett and Professor Plum were in the study at the time of the murder. And Mr. 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 pipe. 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: Miles Davis was a trumpeter.

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 specific 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 Davis's choice of instrument from that fact.

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

- One or more of the premises might be false.
- The conclusion might not follow from the premises.

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

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 it is impossible for the premises to be true whilst the conclusion is false. Consider this example:

Oranges are either fruits or musical instruments.
Oranges are not fruits.
So: Oranges are musical instruments.

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

St. Louis is in the United States.
Beijing is in China.
So: Paris is in France.

The premises and conclusion of this argument are, as a matter of fact, all true. But the argument is invalid. If Paris were to declare independence from the rest of France, then the conclusion would be false, even though both of the premises would remain true. Thus, it is *possible* for the premises of this argument to be true and the conclusion false. The argument is therefore 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 will 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 August 1997, it rained in St. Louis.
> In August 1998, it rained in St. Louis.
> In August 1999, it rained in St. Louis.
> In August 2000, it rained in St. Louis.
> So: It rains every August in St. Louis.

This argument generalizes 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 August 2001, it rained in St. Louis; In August 2002.... But, however many premises of this form we add, the argument will remain invalid. Even if it has rained in St. Louis in every August thus far, it remains *possible* that St. Louis will stay dry next August.

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 will 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. Langston Hughes was either born in Illinois or he was a part of the Harlem Renaissance.
2. Langston Hughes was not a part of the Harlem Renaissance.
So: Langston Hughes 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. Maya Angelou was either from St. Louis or from Chicago.
2. Maya Angelou was not from Chicago.
So: Maya Angelou was from St. Louis.

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. Eliza is now 12 years old.
   2. Eliza is now 3 years old.
So: Oliver 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, we introduced the idea of a valid argument. We will want to introduce some more ideas that are important in logic.

3.1 Truth values

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

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

The common feature of these three kinds of sentence is that they 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 just True and False.

3.2 Consistency

Consider these two sentences:

B1. Etta's only brother is shorter than her.
B2. Etta'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. Willie Nelson recorded "On The Road Again."
2. Someone once recorded "On The Road Again."
3. No one has ever recorded "On The Road Again."
4. If Willie Nelson recorded "On The Road Again," then someone has.
5. Even though Willie Nelson recorded "On The Road Again," no one has ever recorded "On The Road Again."
6. If anyone has ever recorded "On The Road Again," it was Willie Nelson.

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.

Chapter 2

truth-functional logic

First steps to symbolization 4

4.1 Validity in virtue of form

Consider this argument:

> It is sunny outside.
> If it is sunny outside, then Maya is going for a walk.
> So: Maya is going for a walk.

and another argument:

> Nina sings the blues.
> If Nina sings the blues, then Langston will buy a new record player.
> So: Langston will buy a new record player.

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:

> Maya is either going for a walk or at home.
> Maya is not at home.
> So: Maya is going for a walk.

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 Zora both wrote novels and danced competitive tango.

> Zora wrote novels.
> So: Zora did not dance competitive tango.

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 'Maya is going for a walk', 'Langston will buy a new record player', or 'Zora danced competitive tango'. If it has to do with meanings at all, it is with the meanings of phrases like 'and', 'or', 'not,' and 'if..., then...'.

In this chapter, we are going to develop a formal language which allows us to symbolize 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:

> Miles was born in Alton, Illinois.
> So: Miles was born in America.

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:

> Miles was born in Alton, Illinois.
> So: Miles was born in Germany.

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 purple 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 purple 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. (Albert might paint his nails with an elegant shiny purple polish.) Plausibly, the validity of the first argument is keyed to the way that colors (or color-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 *subsentences* of sentences with individual letters. Thus in the first example of this section, 'it is sunny outside' is a subsentence of 'If it is sunny outside, then Maya is going for a walk', 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 will use atomic sentence to represent, or symbolize, certain English sentences. To do this, we provide a SYMBOLIZATION KEY, such as the following:

> A: It is sunny outside.
> C: Maya is going for a walk.

In doing this, we are not fixing this symbolization *once and for all*. We are just saying that, for the time being, we will think of the atomic sentence of TFL, 'A', as symbolizing the English sentence 'It is sunny outside', and the atomic sentence of TFL, 'C', as symbolizing the English sentence 'Maya is going for a walk'. Later, when we are dealing with different sentences or different arguments, we can provide a new symbolization key; as it might be:

> A: Miles is a trumpeter.
> C: Zora dances the tango.

But it is important to understand that whatever structure an English sentence might have is lost when it is symbolized 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 symbolizing 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 will make use of logical connectives to build complex sentences from atomic components. There are five logical connectives in TFL. This table summarizes 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 symbolize these sentences:

1. Willie is in Austin.
2. It is not the case that Willie is in Austin.
3. Willie is not in Austin.

In order to symbolize sentence 1, we will need an atomic sentence. We might offer this symbolization key:

B: Willie is in Austin.

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

15

> A sentence can be symbolized as $\neg\mathcal{A}$ 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 record player can be replaced if it breaks.
5. The record player is irreplaceable.
6. The record player is not irreplaceable.

Let us use the following representation key:

R: The record player is replaceable

Sentence 4 can now be symbolized by 'R'. Moving on to sentence 5: saying the record player is irreplaceable means that it is not the case that the record player is replaceable. So even though sentence 5 does not contain the word 'not', we will symbolize it as follows: '$\neg R$'.

Sentence 6 can be paraphrased as 'It is not the case that the record player is irreplaceable.' Which can again be paraphrased as 'It is not the case that it is not the case that the record player is replaceable'. So we might symbolize this English sentence with the TFL sentence '$\neg\neg R$'. (In English, double-negation tends to cancel out: sentence 6 says something very similar to 'the record player is replaceable'.)

But some care is needed when handling negations. Consider:

7. Maya is happy.
8. Maya is unhappy.

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

5.2 Conjunction

Consider these sentences:

9. Miles is musical.
10. Nina is musical.
11. Miles is musical, and Nina is also musical.

We will need separate atomic sentences of TFL to symbolize sentences 9 and 10; perhaps

A: Miles is musical.
B: Nina is musical.

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

Notice that we make no attempt to symbolize 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) symbolize such things in TFL.

Some more examples will bring out this point:

12. Nina is musical and energetic.
13. Nina and Miles are both musical.
14. Although Nina is musical, she is not energetic.
15. Miles is musical, but Nina is more musical than him.

Sentence 12 is obviously a conjunction. The sentence says two things (about Nina). In English, it is permissible to refer to Nina only once. It *might* be tempting to think that we need to symbolize sentence 12 with something along the lines of 'B and energetic'. This would be a mistake. Once we symbolize 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 Nina is energetic'. So we need to add another sentence letter to the symbolization key. Let 'E' symbolize 'Nina is energetic'. Now the entire sentence can be symbolized as '$(B \land E)$'.

Sentence 13 says one thing about two different subjects. It says of both Nina and Miles that they are musical, and in English we use the word 'musical' only once. The sentence can be paraphrased as 'Nina is musical, and Miles is musical'. We can symbolize this in TFL as '$(B \land A)$', using the same symbolization 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 Nina is musical and that she is not energetic. In order to make each of the conjuncts an atomic sentence, we need to replace 'she' with 'Nina'. So we can paraphrase sentence 14 as, '*Both* Nina is musical, *and* Nina is not energetic'. The second conjunct contains a negation, so we paraphrase further: '*Both* Nina is musical *and it is not the case that* Nina is energetic'. And now we can symbolize this with the TFL sentence '$(E \land \neg B)$'. Note that we have lost all sorts of nuance in this symbolization. There is a distinct difference in tone between sentence 14 and 'Both Nina is musical and it is not the case that Nina is energetic'. 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* Miles is musical, *and* Nina is more musical than Miles'. (Notice that we once again replace the pronoun 'him' with 'Miles'.) How should we deal

with the second conjunct? We already have the sentence letter '*A*', which is being used to symbolize 'Miles is musical', and the sentence '*B*' which is being used to symbolize 'Nina is musical'; but neither of these concerns their relative musicality. So, to to symbolize the entire sentence, we need a new sentence letter. Let the TFL sentence '*R*' symbolize the English sentence 'Nina is more musical than Miles'. Now we can symbolize sentence 15 by '$(A \land R)$'.

> A sentence can be symbolized as $(\mathcal{A} \land \mathcal{B})$ if it can be paraphrased in English as 'Both..., and...', or as '..., but ...', or as 'although ..., ...'.

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

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

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

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

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

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

These English sentences are very different, and their symbolizations 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 Albert will go climbing with me, or he will watch a movie.
19. Either Albert or Oliver will go climbing with me.

For these sentences we can use this symbolization key:

D: Albert will go climbing with me.
E: Oliver will go climbing with me.
M: Albert will watch a movie.

However, we will again need to introduce a new symbol. Sentence 18 is symbolized 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 Albert will go climbing with me, or Oliver will go climbing with me'. Now we can obviously symbolize it by '$(D \vee E)$' again.

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

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

At other times, the word 'or' allows for the possibility that both disjuncts might be true. This is probably the case with sentence 19, above. I might go climbing with Albert, with Oliver, or with both Albert and Oliver. Sentence 19 merely says that I will play with *at least* one of them. This is called an INCLUSIVE OR. The TFL symbol '$\vee$' always symbolizes an *inclusive or*.

It might help to see negation interact 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 symbolization 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 symbolize this in TFL, we need both disjunction and negation. 'It is not the case that you get soup' is symbolized by '$\neg S_1$'. 'It is not the case that you get salad' is symbolized by '$\neg S_2$'. So sentence 20 itself is symbolized 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 symbolize 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 symbolize this as '$(S_1 \vee S_2)$'. The second part says that you do not get both. We can paraphrase this as: 'It is not the case both that you get soup and that you get salad'. Using both negation and conjunction, we symbolize 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 symbolized with '$\wedge$'. Sentence 22 can thus be symbolized as '$((S_1 \vee S_2) \wedge \neg(S_1 \wedge S_2))$'.

   This last example shows something important. Although the TFL symbol '∨' always symbolizes *inclusive or*, we can symbolize an *exclusive or* in TFL. We just have to use a few of our other symbols as well.

5.4   Conditional

Consider these sentences:

  23. If Nina is in Paris, then Nina is in France.
  24. Nina is in France only if Nina is in Paris.

Let's use the following symbolization key:

   P: Nina is in Paris.
   F: Nina is in France

Sentence 23 is roughly of this form: 'if P, then F'. We will use the symbol '→' to symbolize this 'if..., then...' structure. So we symbolize sentence 23 by '$(P \to F)$'. The connective is called THE CONDITIONAL. Here, 'P' is called the ANTECEDENT of the conditional '$(P \to F)$', 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 symbolize 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 Nina to be in Paris that doesn't involve Nina being in France. But sentence 24 is not so straightforward: were Nina in Marsalles, Lyons, or Toulouse, Nina 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 Nina is in France, then Nina is in Paris'. So we can symbolize it by '$(F \to P)$'.

> A sentence can be symbolized as $\mathcal{A} \to \mathcal{B}$ if it can be paraphrased in English as 'If A, then B' or 'A only if B'.

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

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

If we think really hard, all four of these sentences mean the same as 'If Nina is in Paris, then Nina is in France'. So they can all be symbolized by '$P \to F$'.

   It is important to bear in mind that the connective '→' 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 '→' to symbolize English conditionals. We will return to this in §§9.3 and 11.5.

5.5   Biconditional

Consider these sentences:

29. Oliver is a cat only if it he is a mammal.
30. Oliver is a cat if he is a mammal.
31. Oliver is a cat if and only if he is a mammal.

We will use the following symbolization key:

C: Oliver is a cat.
M: Oliver is a mammal.

Sentence 29, for reasons discussed above, can be symbolized by '$C \to M$'.

Sentence 30 is importantly different. It can be paraphrased as, 'If Oliver is a mammal then Oliver is a cat'. So it can be symbolized by '$M \to C$'.

Sentence 31 says something stronger than either 29 or 30. It can be paraphrased as 'Oliver is a cat if he is a mammal, and Oliver is a cat only if Oliver is a mammal'. This is just the conjunction of sentences 29 and 30. So we can symbolize it as '$(C \to M) \land (M \to C)$'. We call this a BICONDITIONAL, because it entails the conditional in both directions.

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

The expression 'if and only if' occurs a lot in philosophy and logic. For brevity, we can abbreviate it with the snappier word 'iff'. I will 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 symbolized as $\mathcal{A} \leftrightarrow \mathcal{B}$ if it can be paraphrased in English as 'A iff B'; that is, as 'A if and only if B'.

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 dessert'. Suppose you ate your greens, but that your parents refused to give you any dessert, on the grounds that they were only committed to the *conditional* (roughly 'if you get dessert, then you will have eaten your greens'), rather than the biconditional (roughly, 'you get dessert 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 you use the biconditional iff you mean to.

5.6   Unless

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

32. Unless you plug up your ears, you will be moved by *A Love Supreme*.
33. You will be moved by *A Love Supreme* unless you plug up your ears.

These two sentences are clearly equivalent. To symbolize them, we will use the symbolization key:

> J: You plug up your ears.
> D: You will be moved by *A Love Supreme*.

Both sentences mean that if you do not plug up your ears, then you will be moved by *A Love Supreme*. With this in mind, we might symbolize them as '$\neg J \to D$'.

Equally, both sentences mean that if you are not moved by *A Love Supreme*, then you have plugged up your ears. With this in mind, we might symbolize them as '$\neg D \to J$'.

Equally, both sentences mean that either you will plug up your ears or you will be moved by *A Love Supreme*. With this in mind, we might symbolize them as '$J \vee D$'.

All three are correct symbolizations. Indeed, in chapter **??** we will see that all three symbolizations are equivalent in TFL.

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

Again, though, there is a little complication. 'Unless' can be symbolized 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 symbolized 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 will go running unless it rains'. I probably mean something like 'I will go running iff it does not rain' (i.e. the biconditional), or 'either I will go running or it will rain, but not both' (i.e. exclusive disjunction). Again: be aware of this when interpreting what other people have said, but be precise in your writing, unless you want to be deliberately ambiguous.

Practice exercises

A. Using the symbolization key given, symbolize 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 symbolization key given, symbolize 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 symbolization key given, symbolize each English sentence in TFL.

    T_1: Miles is a trumpeter.
    T_2: John is a trumpeter.
    S_1: Miles is a saxophonist.
    S_2: John is a saxophonist.
    C_1: Miles has a satisfying career.

C_2: John has a satisfying career.

1. Miles and John are both trumpeters.
2. If Miles is a saxophonist, then he has a satisfying career.
3. Miles is a saxophonist, unless he is an trumpeter.
4. John is an unsatisfied trumpeter.
5. Neither Miles nor John is an trumpeter.
6. Both Miles and John are trumpeters, but neither of them find it satisfying.
7. John is satisfied only if he is a saxophonist.
8. If Miles is not an trumpeter, then neither is John, but if Miles is, then John is too.
9. Miles is satisfied with his career if and only if John is not satisfied with his.
10. If John is both an trumpeter and a saxophonist, then he must be satisfied with his career.
11. It cannot be that John is both an trumpeter and a saxophonist.
12. John and Miles are both saxophonists if and only if neither of them is an trumpeter.

D. Give a symbolization key and symbolize the following sentences in TFL.

1. Elizabeth and Phillip are both spies.
2. If either Elizabeth or Phillip is a spy, then the code has been broken.
3. If neither Elizabeth nor Phillip is a spy, then the code remains unbroken.
4. The Russian 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 Russian embassy will be in an uproar regardless.
6. Either Elizabeth or Phillip is a spy, but not both.

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

1. If there is work to be done by the Free French Resistance, then Josephine Baker will perform.
2. Josephine Baker will not perform unless she is supplied with a visa.
3. Josephine Baker will either perform or she won't, but there is work to be done by the Free French Resistance regardless.
4. Charles de Gaulle will lead the Free French Resistance if and only if Germany continues to occupy France.
5. If Charles de Gaulle will lead the Free French Resistance, Josephine Baker will be supplied with a visa.

F. For each argument, write a symbolization key and symbolize all of the sentences of the argument in TFL.

1. If Etta plays the piano in the morning, then Artis wakes up cranky. Etta plays piano in the morning unless she is distracted. So if Artis does not wake up cranky, then Etta must be distracted.

2. It will either rain or snow on Tuesday. If it rains, Dorothy will be sad. If it snows, Dorothy will be cold. Therefore, Dorothy will either be sad or cold on Tuesday.
3. If Langston takes Zora out on the town, then things are thrilling, but not predictable. If he doesn't, then things are predictable, but not exciting. Therefore, things are either exciting or predictable; but not both.

G. *optional* We symbolized an *exclusive or* using '∨', '∧', and '¬'. How could you symbolize an *exclusive or* using only two connectives? Is there any way to symbolize 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 '$(A \lor B)$' 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, we will offer a complete *definition* of what counts 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	$A, B, C, \ldots, Z$
with subscripts, as needed	$A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \ldots$
Connectives	$\neg, \land, \lor, \rightarrow, \leftrightarrow$
Brackets	(,)

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

6.2 Sentences

Of course, many expressions of TFL will be total gibberish. We want to know when an expression of TFL amounts to a *sentence.*

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 '$(A \land G_{13})$', '$(G_{13} \land A)$', '$(A \land A)$', and '$(G_{13} \land G_{13})$'. We could also apply negation repeatedly to get sentences like '$\neg\neg A$' or apply negation along with conjunction to get sentences like '$\neg(A \land G_{13})$' and '$\neg(G_{13} \land \neg G_{13})$'. The possible combinations are endless, even starting with just these two sentence letters, and there are infinitely many sentence letters. So there is no point in trying to list all the sentences one by one.

Instead, we will describe the process by which sentences can be *constructed*. Consider negation: Given any sentence $\mathcal{A}$ of TFL, $\neg\mathcal{A}$ 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 $\mathcal{A}$ and $\mathcal{B}$ are sentences of TFL, then $(\mathcal{A} \wedge \mathcal{B})$ 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 $\mathcal{A}$ is a sentence, then $\neg\mathcal{A}$ is a sentence.

3. If $\mathcal{A}$ and $\mathcal{B}$ are sentences, then $(\mathcal{A} \wedge \mathcal{B})$ is a sentence.

4. If $\mathcal{A}$ and $\mathcal{B}$ are sentences, then $(\mathcal{A} \vee \mathcal{B})$ is a sentence.

5. If $\mathcal{A}$ and $\mathcal{B}$ are sentences, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a sentence.

6. If $\mathcal{A}$ and $\mathcal{B}$ are sentences, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a sentence.

7. Nothing else is a sentence.

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

- My parents are ancestors of mine.

and then offer further clauses like:

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

Using this definition, we can easily check to see whether someone is my ancestor: just check whether she is the parent of the parent of. . . one of my parents. And the same is true for our recursive definition of sentences of 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 get down to atomic sentences, then we are ok.

Let's consider some examples.

Suppose we want to know whether or not '$\neg\neg\neg D$' is a sentence of 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 *last*, 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 '$((\neg E \vee F) \to \neg\neg G)$', the main logical operator is '$\to$'.

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. We will return to this point in chapter **??**.

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. So in a sentence like:

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

this was constructed by conjoining 'P' with '$(\neg(R \wedge B) \leftrightarrow Q)$'. This last sentence was constructed by placing a biconditional between '$\neg(R \wedge B)$' and 'Q'. And the former of these sentences—a subsentence of our original sentence—is a sentence for which '$\neg$' is the main 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 '$Q \land R$' instead of the sentence '$(Q \land R)$'. However, we must remember to put the brackets back in, when we want to embed the sentence into a more complicated sentence!

Second, it can be a bit painful to stare at long sentences with many nested pairs of brackets. To make things a bit easier on the eyes, we will allow ourselves to use square brackets, '[' and ']', instead of rounded ones. So there is no logical difference between '$(P \lor Q)$' and '$[P \lor Q]$', for example.

Combining these two conventions, we can rewrite the unwieldy sentence

$$(((H \to I) \lor (I \to H)) \land (J \lor K))$$

rather more simply as follows:

$$\big[(H \to I) \lor (I \to H)\big] \land (J \lor K)$$

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. $J_{374} \lor \neg J_{374}$
3. $\neg\neg\neg\neg F$
4. $\neg \land S$
5. $(G \land \neg G)$
6. $(A \to (A \land \neg F)) \lor (D \leftrightarrow E)$
7. $[(Z \leftrightarrow S) \to W] \land [J \lor X]$
8. $(F \leftrightarrow \neg D \to J) \lor (C \land D)$

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?

$$\big[(H \to I) \lor (I \to H)\big] \land (J \lor K)$$

Use and mention   7

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

7.1 Quotation conventions

Consider these two sentences:

- Lyda Krewson is the Mayor of St. Louis.
- The expression 'Lyda Krewson' is composed of two uppercase letters and nine lowercase letters

When we want to talk about the Mayor, we *use* her name. When we want to talk about the Mayor's name, we *mention* that name. And we do so by putting it in quotation marks.

There is a general point here. When we want to talk about things in the world, we just *use* words. When we want to talk about words, we typically have to *mention* those words. 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:

- 'Lyda Krewson' is the Mayor of St. Louis.

says that some *expression* is the Mayor. And that's false. The *person* is the Mayor; her *name* isn't. Conversely, this sentence:

- Lyda Krewson is composed of two uppercase letters and ten lowercase letters.

also says something false: Lyda Krewson is a person, made of flesh and bones rather than letters. One final example:

- " 'Lyda Krewson' " is the name of 'Lyda Krewson'.

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

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

7.2 Object language and metalanguage

These general quotation conventions are of particular importance for us. After all, we are describing a formal language here, TFL, and so we are often *mentioning* 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 exactly, but English supplemented with some additional vocabulary which helps us to get along.

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

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

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   Swash-fonts and Quine quotes

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 uppercase swash-font letters to do this, namely:

$$\mathcal{A}, \mathcal{B}, \mathcal{C}, \mathcal{D}, \ldots$$

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 repeat the second clause of the recursive definition of a sentence of TFL, we said:

   3. If $\mathcal{A}$ is a sentence, then $\neg\mathcal{A}$ 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 emphasize, then:

> '$\mathcal{A}$' is a symbol in augmented English, which we use to talk about any TFL expression. 'A' is a particular atomic sentence of TFL.

But this last example raises a further complications for our quotation conventions. I have not included 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 this rule is not a sentence of English, since it contains '$\neg$'. So we might try to write:

   3′. If $\mathcal{A}$ is a sentence, then '$\neg\mathcal{A}$' is a sentence.

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

   What we really want to say is something like this:

   3″. If $\mathcal{A}$ and $\mathcal{B}$ are sentences, then the result of concatenating the symbol '$\neg$' with the sentence $\mathcal{A}$ is a sentence.

This is impeccable, but rather long-winded. But we can avoid long-windedness by creating our own conventions. We can perfectly well stipulate that an expression like '$\neg\mathcal{A}$' should simply be read *directly* in terms of rules for concatenation. So, *officially*, the metalanguage expression '$\neg\mathcal{A}$' simply abbreviates:

   the result of concatenating the symbol '$\neg$' with the sentence $\mathcal{A}$

and similarly, for expressions like '$(\mathcal{A} \wedge \mathcal{B})$', '$(\mathcal{A} \vee \mathcal{B})$', 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 symbolize English sentences, we can symbolize English arguments using TFL. Thus we might ask whether the argument whose premises are the TFL sentences 'A' and '$A \rightarrow C$', and whose conclusion is the TFL sentence 'C', is valid. However, it is quite a mouthful to write that every time. So instead I will introduce another bit of abbreviation. This:

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

abbreviates:

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

To avoid unnecessary clutter, we will not regard this as requiring quotation marks around it. (Note, then, that '$\therefore$' is a symbol of our augmented *metalanguage*, and not a new symbol of TFL.)

Chapter 3

truth tables

Characteristic truth tables   8

Any non-atomic sentence of TFL is composed of atomic sentences with senten-
tial 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 '$(D \land E)$', for instance, you only need to know the truth value
of 'D' and the truth value of 'E'.

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

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

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

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

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

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

Disjunction.   Recall that '$\lor$' always represents inclusive or. So, for any sen-
tences $\mathcal{A}$ and $\mathcal{B}$, $\mathcal{A} \lor \mathcal{B}$ is true if and only if either $\mathcal{A}$ or $\mathcal{B}$ is true. We can
summarize this in the characteristic truth table for disjunction:

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

Like conjunction, disjunction is symmetrical.

Conditional.   I'm just going to come clean and admit it. Conditionals are a right old mess in TFL. Exactly how much of a mess they are is a matter of *philosophical* contention. We'll look more into these subtleties in §§9.3 and 11.5. For now, let's stipulate the following: $\mathcal{A} \to \mathcal{B}$ is false if and only if $\mathcal{A}$ is true and $\mathcal{B}$ is false. We can summarize this with a characteristic truth table for the conditional.

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

The conditional is *asymmetrical*. You cannot swap the antecedent and consequent without changing the meaning of the sentence, because $\mathcal{A} \to \mathcal{B}$ has a very different truth table from $\mathcal{B} \to \mathcal{A}$.

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

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

Unsurprisingly, the biconditional is symmetrical.

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*.

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

1. $2 + 2 = 4$
2. John Coltrane wrote "A Love Supreme"

Whereas it is necessarily the case that $2 + 2 = 4$, it is not *necessarily* the case that John Coltrane wrote "A Love Supreme". If Coltrane had died earlier, he might never have written it. So 'It is necessarily the case that...' is a connective of English, but it is not *truth-functional*.

9.2 Symbolizing 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 symbolize 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. Eliza is tiny and Eliza is mighty.
2. Although Eliza is tiny, Eliza is mighty.
3. Despite being tiny, Eliza is mighty.
4. Eliza is mighty, albeit tiny.
5. Eliza's tininess notwithstanding, he is mighty.

All of the above sentences will be symbolized with the same TFL sentence, perhaps '$F \wedge Q$'.

I keep saying that we use TFL sentences to *symbolize* 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 just cannot do that. This is why we're talking about *symbolizing* English sentences, rather than of *translating* them.

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

F: Eliza is tiny.
Q: Eliza is mighty.

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

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

9.3   Indicative versus subjunctive conditionals

I want to bring home the point that TFL can *only* deal with truth functions by considering the case of the conditional. When we introduced the characteristic truth table for the material conditional in §8, we didn't do anything to *justify* it. But now we're in a position to offer a justification, which follows Dorothy Edgington.[1]

---

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

Suppose that Dorothy has drawn some shapes on a piece of paper, and colored 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, Dorothy has drawn the following:

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

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

However, if Dorothy had drawn a fourth shape, thus:
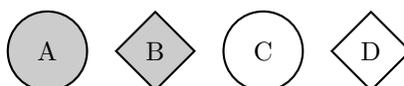
then my claim would have be false. So it must be that this claim is false:

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

Sentence 1 is true; sentence 2 is false. But both have false antecedents and false consequents. So the truth value of the whole sentence is not uniquely determined by the truth value of the parts. Do not just blithely assume that you can adequately symbolize an English 'if..., then...' with TFL's '→'.

The crucial point is that sentences 1 and 2 employ *subjunctive* conditionals, rather than *indicative* conditionals. They ask us to imagine something contrary to fact—Mitt Romney lost the 2012 election—and then ask us to evaluate what *would* have happened in that case. Such considerations just cannot be tackled using '→'.

We'll look into more difficulties raised by conditionals in §11.5. For now, we'll need to be content with the observation that '→' is the only candidate for a truth-functional conditional, but that many English conditionals cannot be represented adequately using '→'. TFL is an intrinsically limited language.

Complete truth tables   10

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

> A VALUATION is any assignment of truth values to 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 entire truth table represents all possible valuations. And the truth table provides us with a means to calculate the truth value of complex sentences, on each possible valuation. This is easiest to explain by example.

10.1 A worked example

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

H	I	$(H \land I) \to H$
T	T	
T	F	
F	T	
F	F	

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

H	I	$(H$	$\land I)$	$\to H$
T	T	T	T	T
T	F	T	F	T
F	T	F	T	F
F	F	F	F	F

41

Now consider the subsentence '$(H \wedge I)$'. This is a conjunction, $(\mathcal{A} \wedge \mathcal{B})$, with '$H$' as $\mathcal{A}$ and with 'I' as $\mathcal{B}$. The characteristic truth table for conjunction gives the truth conditions for *any* sentence of the form $(\mathcal{A} \wedge \mathcal{B})$, whatever $\mathcal{A}$ and $\mathcal{B}$ might be. It summarizes 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.

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

Now, the entire sentence that we are dealing with is a conditional, $\mathcal{A} \to \mathcal{B}$, with '$(H \wedge I)$' as $\mathcal{A}$ and with 'H' as $\mathcal{B}$. On the second row, for example, '$(H \wedge I)$' is false and 'H' is true. Since a conditional is true when the antecedent is false, we write a 'T' in the second row underneath the conditional symbol. We continue for the other three rows and get this:

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

The conditional is the main logical connective of the sentence. And the column of 'T's underneath the conditional tells us that the sentence '$(H \wedge I) \to H$' 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 '$(H \wedge I) \to H$' is true on every valuation.

In this example, you can see that that 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	$(H \wedge I) \to H$		
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. This is emphasized by putting this column in bold. When you work through truth tables yourself, you should similarly emphasize it by underlining or circling the relevant column.

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 '$[(C \leftrightarrow C) \to C] \wedge \neg (C \to C)$'. 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	$[(C \leftrightarrow C) \to C] \wedge \neg (C \to C)$
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.

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

A sentence that contains three atomic sentences requires eight lines:

M	N	P	$M \wedge (N \vee P)$
T	T	T	T **T** T T T
T	T	F	T **T** T T F
T	F	T	T **T** F T T
T	F	F	T **F** F F F
F	T	T	F **F** T T T
F	T	F	F **F** T T F
F	F	T	F **F** F T T
F	F	F	F **F** F F F

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

A complete truth table for a sentence that contains four different atomic sentences requires 16 lines. Five letters, 32 lines. Six letters, 64 lines. And so

on. To be perfectly general: If a complete truth table has n different atomic sentences, then it must have 2^n lines.

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, we're allowed 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:

$$((A \vee B) \vee C)$$
$$(A \vee (B \vee C))$$

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:

$$((A \rightarrow B) \rightarrow C)$$
$$(A \rightarrow (B \rightarrow C))$$

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:

$$((A \lor B) \land C)$$
$$(A \lor (B \land C))$$

So if we were to write:

$$A \lor B \land C$$

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. $A \to A$
2. $C \to \neg C$
3. $(A \leftrightarrow B) \leftrightarrow \neg(A \leftrightarrow \neg B)$
4. $(A \to B) \lor (B \to A)$
5. $(A \land B) \to (B \lor A)$
6. $\neg(A \lor B) \leftrightarrow (\neg A \land \neg B)$
7. $\big[(A \land B) \land \neg(A \land B)\big] \land C$
8. $\big[(A \land B) \land C\big] \to B$
9. $\neg\big[(C \lor A) \lor B\big]$

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

1. '$((A \land B) \land C)$' and '$(A \land (B \land C))$' have the same truth table
2. '$((A \lor B) \lor C)$' and '$(A \lor (B \lor C))$' have the same truth table
3. '$((A \lor B) \land C)$' and '$(A \lor (B \land C))$' do not have the same truth table
4. '$((A \to B) \to C)$' and '$(A \to (B \to C))$' 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, we 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, we will introduce some related ideas, and show how to use truth tables to test whether or not they apply.

11.1 Tautologies and contradictions

In §3.1, we learned about *necessary truth* and *necessary falsity*. Both notions have surrogates in TFL. We will start with a surrogate for necessary truth.

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

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

This is only, though, a surrogate for necessary truth. There are some necessary truths that we cannot adequately symbolize in TFL. An example is '$2 + 2 = 4$'. This *must* be true, but if we try to symbolize it in TFL, the best we can offer is an atomic sentence, and no atomic sentence is a tautology. Still, if we can adequately symbolize 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:

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

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

11.2 Tautological equivalence

Here is a similar, useful notion:

> $\mathcal{A}$ and $\mathcal{B}$ 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 '$(A \land B) \land C$' and '$A \land (B \land C)$' are tautologically equivalent. Again, it is easy to test for tautological equivalence using truth tables. Consider the sentences '$\neg(P \lor Q)$' and '$\neg P \land \neg Q$'. Are they tautologically equivalent? To find out, we construct a truth table.

P	Q	$\neg\ (P \lor Q)$	$\neg P \land \neg Q$
T	T	**F** T T T	F T **F** F T
T	F	**F** T T F	F T **F** T F
F	T	**F** F T T	T F **F** F T
F	F	**T** F F F	T F **T** T F

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

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

Derivatively, sentences are jointly tautologically inconsistent if there is no valuation that 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:

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

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

J	L	$\neg L \to (J \vee L)$	$\neg\ L$	J
T	T	F T **T** T T T	**F** T	**T**
T	F	T F **T** T T F	**T** F	**T**
F	T	F T **T** F T T	**F** T	**F**
F	F	T F **F** F F F	**T** F	**F**

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

We now make an important observation:

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

Here's why. If $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ tautologically entail $\mathcal{C}$, then there is no valuation which makes all of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ true whilst making $\mathcal{C}$ false. This means that it is *logically impossible* for $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ all to be true whilst $\mathcal{C}$ is false. But this is just what it takes for an argument, with premises $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ and conclusion $\mathcal{C}$, to be valid!

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

11.5   The limits of these tests

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

First, consider the argument:

1. Eliza has four legs. So Eliza has more than two legs.

To symbolize 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 surely seems valid!

Second, consider the sentence:

2. Willie is neither bald nor not-bald.

To symbolize 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 go on to add 'Willie is on the borderline of baldness'!

Third, consider the following sentence:

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

Symbolizing this in TFL, we would offer something like '$\neg(G \to M)$'. Now, '$\neg(G \to M)$' tautologically entails 'G' (again, check this with a truth table). So if we symbolize 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 Willie'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

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

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

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

But let me be clear. '$\vDash$' 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:

- $P, P \to Q \vDash Q$

is just an abbreviation for the English sentence:

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

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

$$\vDash \mathcal{C}$$

This says that there is no valuation which makes all the sentences mentioned on the left side of '$\vDash$' true whilst making $\mathcal{C}$ false. Since *no* sentences are mentioned on the left side of '$\vDash$' in this case, this just means that there is no valuation which makes $\mathcal{C}$ false. Otherwise put, it says that every valuation makes $\mathcal{C}$ true. Otherwise put, it says that $\mathcal{C}$ is a tautology. Equally:

$$\mathcal{A} \vDash$$

says that $\mathcal{A}$ is a contradiction.

11.7 '⊨' versus '→'

I now want to compare and contrast '⊨' and '→'.

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

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

Combining these two observations, we see that $\mathcal{A} \to \mathcal{C}$ is a tautology iff $\mathcal{A} \vDash \mathcal{C}$. But there is a really, really important difference between '⊨' and '→':

> '→' is a sentential connective of TFL.
> '⊨' is a symbol of augmented English.

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

Practice exercises

A. Revisit your answers to §10**A**. 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. $A \to A$, $\neg A \to \neg A$, $A \wedge A$, $A \vee A$
2. $A \vee B$, $A \to C$, $B \to C$
3. $B \wedge (C \vee A)$, $A \to B$, $\neg(B \vee C)$
4. $A \leftrightarrow (B \vee C)$, $C \to \neg A$, $A \to \neg B$

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

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

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

1. Suppose that $\mathcal{A}$ and $\mathcal{B}$ are tautologically equivalent. What can you say about $\mathcal{A} \leftrightarrow \mathcal{B}$?
2. Suppose that $(\mathcal{A} \wedge \mathcal{B}) \to \mathcal{C}$ is neither a tautology nor a contradiction. What can you say about whether $\mathcal{A}, \mathcal{B}$ ∴ $\mathcal{C}$ is valid?

3. Suppose that $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ are jointly tautologically inconsistent. What can you say about $(\mathcal{A} \land \mathcal{B} \land \mathcal{C})$?
4. Suppose that $\mathcal{A}$ is a contradiction. What can you say about whether $\mathcal{A}, \mathcal{B} \vDash \mathcal{C}$?
5. Suppose that $\mathcal{C}$ is a tautology. What can you say about whether $\mathcal{A}, \mathcal{B} \vDash \mathcal{C}$?
6. Suppose that $\mathcal{A}$ and $\mathcal{B}$ are tautologically equivalent. What can you say about $(\mathcal{A} \lor \mathcal{B})$?
7. Suppose that $\mathcal{A}$ and $\mathcal{B}$ are *not* tautologically equivalent. What can you say about $(\mathcal{A} \lor \mathcal{B})$?

E. Consider the following principle:

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

Is this principle correct? Explain your answer.

Truth table shortcuts 12

With practice, you will quickly become adept at filling out truth tables. In this section, we'll see some permissible shortcuts to help us 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	**F**	F
T	F	T	**F**	F
F	T	T	**T**	T
F	F	F	**F**	T

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	F	F	**F**	F
T	F	F	T	T	**T**	F
F	T				**T**	T
F	F				**T**	T

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				**F**	F
T	F				**F**	F
F	T	T	F		**T**	T
F	F	T	F		**T**	T

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	$((P \to Q) \to P) \to P$		
T	T			**T**
T	F			**T**
F	T	T	F	**T**
F	F	T	F	**T**

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

12.2   Testing for validity and entailment

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

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

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

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

$$\neg L \to (J \lor L), \neg L \therefore J$$

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

J	L	$\neg L \to (J \lor L)$	$\neg L$	J
T	T			T
T	F			T
F	T	?	?	F
F	F	?	?	F

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

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

J	L	$\neg L \to (J \lor L)$	$\neg L$	J
T	T			T
T	F			T
F	T		F	F
F	F	?	T	F

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

J	L	¬L → (J ∨ L)			¬L	J
T	T					T
T	F					T
F	T				F	F
F	F	T	**F**	F	T	F

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

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

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

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

A	B	C	D	A ∨ B	¬(A ∧ C)	¬(B ∧ ¬D)	(¬C ∨ D)	
T	T	T	T					**T**
T	T	T	F	?	?	?	F	**F**
T	T	F	T					**T**
T	T	F	F				T	**T**
T	F	T	T					**T**
T	F	T	F	?	?	?	F	**F**
T	F	F	T					**T**
T	F	F	F				T	**T**
F	T	T	T					**T**
F	T	T	F	?	?	?	F	**F**
F	T	F	T					**T**
F	T	F	F				T	**T**
F	F	T	T					**T**
F	F	T	F	?	?	?	F	**F**
F	F	F	T					**T**
F	F	F	F				T	**T**

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

A	B	C	D	A ∨ B	¬ (A∧C)		¬ (B∧¬D)		(¬C ∨ D)	
T	T	T	T							**T**
T	T	T	F	**T**	**F**	T			F	**F**
T	T	F	T							**T**
T	T	F	F						T	**T**
T	F	T	T							**T**
T	F	T	F	**T**	**F**	T			F	**F**
T	F	F	T							**T**
T	F	F	F						T	**T**
F	T	T	T							**T**
F	T	T	F	**T**	**T**	F	**F**	TT	F	**F**
F	T	F	T							**T**
F	T	F	F						T	**T**
F	F	T	T							**T**
F	F	T	F	**F**					F	**F**
F	F	F	T							**T**
F	F	F	F						T	**T**

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

I have been discussing shortcuts in testing for tautological 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, determine whether each sentence is a tautology, a contradiction, or neither.

1. ¬B ∧ B
2. ¬D ∨ D
3. (A ∧ B) ∨ (B ∧ A)
4. ¬[A → (B → A)]
5. A ↔ [A → (B ∧ ¬B)]
6. ¬(A ∧ B) ↔ A
7. A → (B ∨ C)
8. (A ∧ ¬A) → (B ∨ C)
9. (B ∧ D) ↔ [A ↔ (A ∨ C)]

Partial truth tables 13

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 '$(U \land T) \to (S \land W)$' is *not* a tautology. We set up a PARTIAL TRUTH TABLE:

S	T	U	W	$(U \land T) \to (S \land W)$
				F

We have only left space for one line, rather than 16, since we are only looking for one line, on which the sentence is false. For just that reason, we have filled in 'F' for the entire sentence.

The main 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	$(U \land T) \to (S \land W)$
				T **F** F

In order for the '$(U \land T)$' to be true, both 'U' and 'T' must be true.

S	T	U	W	$(U \land T) \to (S \land W)$
	T	T		T T T **F** F

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

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

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

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

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

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

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

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. 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 summarizes 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. $A \vee \big[A \rightarrow (A \leftrightarrow A)\big] \therefore A$
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. $A \leftrightarrow B, B \leftrightarrow C \therefore A \leftrightarrow C$

Chapter 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:

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

To symbolize it in TFL, we might offer a symbolization key:

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

And the argument itself becomes:

$$L, A \therefore F$$

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

The problem is not that we have made a mistake while symbolizing the argument. This is the best symbolization 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 symbolization, we lose the connection between Dorothy's being a logician and Dorothy's wearing a hat.

The basic units of TFL are atomic sentences, and TFL cannot decompose these. To symbolize 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 Oliver, or let 'i' stand for Dorothy.

Second, we have predicates. English predicates are expressions like '_____ is a cat' 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 'Oliver is a cat' or 'Dorothy is a logician'. In FOL, we indicate predicates with uppercase italic letters. For instance, we might

let the FOL predicate 'D' symbolize the English predicate '_____ is a cat'. Then the expression 'Db' will be a sentence in FOL, which symbolizes the English sentence 'Oliver is a cat'. Equally, we might let the FOL predicate 'L' symbolize the English predicate '_____ is a logician'. Then the expression 'Li' will symbolize the English sentence 'Dorothy is a logician'.

Third, we have quantifiers. For instance, '$\exists$' will roughly convey 'There is at least one . . .'. So we might symbolize the English sentence 'there is a cat' with the FOL sentence '$\exists x Dx$', which we would naturally read out-loud as 'there is at least one thing, x, such that x is a cat'.

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 'cat' is not a singular term, because there are a great many cats. The phrase 'Eliza' is a singular term, because it refers to a specific cat. Likewise, the phrase 'Kathryn's cat Eliza' is a singular term, because it refers to a specific little cat.

Proper names are a particularly important kind of singular term. These are expressions that pick out individuals without describing them. The name 'Dorothy' is a proper name, and the name alone does not tell you anything about Dorothy. 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. Hilary Putnam, the famous philosopher, is a man. 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:

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

These should be thought of along the lines of proper names in English. But with one difference. 'Kathryn Lindeman' is a proper name, but there are several people with this name. We live with this kind of ambiguity in English, allowing context to individuate the fact that 'Kathryn Lindeman' refers to the professor of this course, and not some other person. 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 symbolization keys. These indicate, temporarily, what a name shall pick out. So we might offer:

   b: Beyoncé
   l: Willie
   e: Eliza

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 cat
   _____ is a member of the Highwaymen
   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, 'Willie borrowed the family car from Beyoncé.' By removing a singular term, we can obtain any of three different predicates:

   _____ borrowed the family car from Beyoncé
   Willie borrowed _____ from Beyoncé
   Willie borrowed the family car from _____

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

   Ex: _____$_x$ is excited
   Nx: _____$_x$ is napping

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

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

1. Beyoncé is excited.
2. Willie and Eliza are excited.
3. If Beyoncé is excited, then so are Willie and Eliza.

Sentence 1 is straightforward: we symbolize it by 'Eb'.

   Sentence 2: this is a conjunction of two simpler sentences. The simple sentences can be symbolized just by 'El' and 'Ee'. Then we help ourselves to our resources from TFL, and symbolize the entire sentence by '$El \wedge Ee$'. 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 symbolize this with '$Eb \rightarrow (El \wedge Ee)$'.

14.4   Quantifiers

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

4. Everyone is excited.
5. Someone is napping.

It might be tempting to symbolize sentence 4 as '$Eb \wedge El \wedge Ee$'. Yet this would only say that Beyoncé, Willie, and Eliza 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 symbolize sentence 4 as '$\forall x Hx$'. 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 '$\forall x Hx$', '$\forall y Hy$', '$\forall z Hz$', and '$\forall x_5 Hx_5$' use different variables, but they will all be logically equivalent.

To symbolize 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 symbolized by '$\exists x Ax$'. Whereas '$\forall x Ax$' is read naturally as 'for all x, x is angry', '$\exists x Ax$' 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 symbolized sentence 5 with '$\exists z Az$', '$\exists w_{256} Aw_{256}$', or whatever.

Some more examples will help. Consider these further sentences:

6. No one is napping.
7. There is someone who is not excited.
8. Not everyone is napping.

Sentence 6 can be paraphrased as, 'It is not the case that someone is napping'. We can then symbolize it using negation and an existential quantifier: '$\neg \exists x Ax$'. Yet sentence 6 could also be paraphrased as, 'Everyone is not napping'. With this in mind, it can be symbolized using negation and a universal quantifier: '$\forall x \neg Ax$'. Both of these are acceptable symbolizations. Indeed, it will transpire that, in general, $\forall x \neg \mathcal{A}$ is logically equivalent to $\neg \exists x \mathcal{A}$. (Notice that I have here returned to the practice of using '$\mathcal{A}$' as a metavariable, from §7.) Symbolizing 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 excited'. This then becomes '$\exists x \neg Ex$'. Of course, we could equally have written '$\neg \forall x Ex$', which we would naturally read as 'it is not the case that everyone is excited'. And that would be a perfectly adequate symbolization of sentence 8.

14.5   Domains

Given the symbolization key we have been using, '$\forall x Ex$' symbolizes 'Everyone is excited'. 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 St. Louis, we define the domain to be people in St. Louis. We write this at the beginning of the symbolization key, like this:

domain: people in St. Louis

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

In FOL, the domain must always include at least one thing. Moreover, in English we can infer 'something is excited' from 'Willie is excited'. In FOL, then, we shall want to be able to infer '$\exists x E x$' from 'El'. 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 St. Louis, then we need to include those people in the domain.

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

Sentences with one quantifier   15

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

15.1   Dealing with syncategorematic adjectives

When we encounter a sentence like

1. Oliver is a black and white cat.

We can paraphrase this as 'Oliver is black and white and Oliver is a cat'. We can then use a symbolization key like:

Bx: _____$_x$ is black and white.
Cx: _____$_x$ is a cat.
    o: Oliver

This allows us to symbolize sentence 1 as '$Bo \land Co$'. 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 Oliver, we might try to use a symbolization key like:

Sx: _____$_x$ is short
Bx: _____$_x$ is a basketball player
Mx: _____$_x$ is a man
    d: Damon Stoudamire

Then we would symbolize sentence 2 with '$Sd \land Bd$', sentence 3 with 'Md' and sentence 4 with '$Sd \land Md$'. 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 symbolize '_____ is

a short basketball player' and '_____ is a short man' using completely different predicates.

Similar examples abound. All musicians are people, but some good musicians are bad people. I might be an incompetent cat burglar, but a competent professor. 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 symbolization 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 need to deal with any names. So here is our key:

domain: all coins
Px: _____$_x$ is in my pocket
Tx: _____$_x$ is on the table
Qx: _____$_x$ is a quarter
Dx: _____$_x$ is a dime

Sentence 5 is most naturally symbolized 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 symbolize 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 symbolize 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 symbolized as $\forall x(\mathcal{F}x \rightarrow \mathcal{G}x)$ if it can be paraphrased in English as 'every F is G'.

Sentence 6 is most naturally symbolized 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 symbolize 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 '$\exists x(Tx \to Dx)$'. That would mean that there is some object in the domain of which '$(Tx \to Dx)$' is true. Recall that, in TFL, $\mathcal{A} \to \mathcal{B}$ is tautologically equivalent to $\neg\mathcal{A} \lor \mathcal{B}$. This equivalence will also hold in FOL. So '$\exists x(Tx \to Dx)$' is true if there is some object in the domain, such that '$(\neg Tx \lor Dx)$' is true of that object. That is, '$\exists x(Tx \to Dx)$' 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 '$\exists x(Tx \to Dx)$' 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 symbolized as $\exists x(\mathcal{F}x \land \mathcal{G}x)$ 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 symbolize it by '$\neg\forall x(Tx \to Dx)$'. You might look at sentence 7 and paraphrase it instead as, 'Some coin on the table is not a dime'. You would then symbolize it by '$\exists x(Tx \land \neg Dx)$'. Although it is probably not immediately obvious yet, these two sentences are logically equivalent. (This is due to the logical equivalence between $\neg\forall x\mathcal{A}$ and $\exists x\neg\mathcal{A}$, mentioned in §14, along with the equivalence between $\neg(\mathcal{A} \to \mathcal{B})$ and $\mathcal{A} \land \neg\mathcal{B}$.)

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

15.3   Empty predicates

In §14, I emphasized 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 symbolize these two sentences:

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

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

domain: animals
  Mx: _____$_x$ is a monkey.
   Sx: _____$_x$ knows sign language.

Sentence 9 can now be symbolized by '$\forall x(Mx \rightarrow Sx)$'. Sentence 10 can be symbolized as '$\exists x(Mx \land Sx)$'.

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

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

Rx: _____x is a refrigerator

Now consider the sentence '$\forall x(Rx \rightarrow Mx)$'. This symbolizes 'every refrigerator is a monkey'. And this sentence is true, given our symbolization key. This is counterintuitive, since we do not want to say that there are a whole bunch of refrigerator monkeys. It is important to remember, though, that '$\forall x(Rx \rightarrow Mx)$' is true iff any member of the domain that is a refrigerator is a monkey. Since the domain is *animals*, there are no refrigerators in the domain. Again, then, the sentence is *vacuously* true.

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

> When $\mathcal{F}$ is an empty predicate, a sentence $\forall x(\mathcal{F}x \rightarrow \ldots)$ will be vacuously true.

15.4 Picking a domain

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

11. Every rose has a thorn.

We might offer this symbolization key:

Rx: _____x is a rose
Tx: _____x has a thorn

It is tempting to say that sentence 11 should be symbolized 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 symbolization. Yet if the domain is merely *things on my*

kitchen table, then '$\forall x(Rx \to Tx)$' 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 symbolize 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 symbolized with '$\forall x Tx$'. 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 symbolize 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; whatevers. And we will certainly need to include a more expansive domain if we simultaneously want to symbolize sentences like:

12. Every cowboy sings a sad, sad song.

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

domain: people and plants
    Cx: _____$_x$ is a cowboy
    Sx: _____$_x$ sings a sad, sad song
    Rx: _____$_x$ is a rose
    Tx: _____$_x$ has a thorn

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

In general, the universal quantifier can be used to symbolize 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 symbolizing English sentences in FOL, it is important to understand the structure of the sentences you want to symbolize. What matters is the final symbolization 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 symbolize directly in FOL.

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

domain: those who can hear
    Mx: _____$_x$ is moved by A Love Supreme.
    Hx: _____$_x$ hears A Love Supreme.
      k: Kathryn

Now consider these sentences:

13. If Kathryn hears A Love Supreme, she will be moved.
14. If a person hears A Love Supreme, then she will be moved.

The same words appear as the consequent in sentences 13 and 14 ('... she will be moved'), 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 Kathryn hears A Love Supreme, then *Kathryn* will be moved'. This can obviously be symbolized as '$Hk \rightarrow Mk$'.

Sentence 14 must be paraphrased differently: 'If a person hears A Love Supreme, then *that person* will be moved'. This sentence is not about any particular person, so we need a variable. As a first pass, we can paraphrase this as, 'For any person x, if x hears A Love Supreme, then x will be moved'. Now this can be symbolized as '$\forall x(Hx \rightarrow Mx)$'. This is the same sentence we would have used to symbolize 'Everyone who hears A Love Supreme is moved'. And on reflection, that is surely true iff sentence 14 is true, as we would hope.

Consider these further sentences:

15. If anyone is unmoved by A Love Supreme, then Kathryn will eat her hat.
16. If anyone is unmoved by A Love Supreme, then they lack taste in music. The same words appear as the antecedent in sentences 15 and 16 ('If anyone is unmoved ...'). But it can be tricky to work out how to symbolize these two uses. Again, paraphrase will come to our aid. Let's add a new predicate to the symbolization key, namely:

    Ex: _____$_x$ will eat their hat.
    Tx: _____$_x$ lacks taste in music.

Sentence 15 can be paraphrased, 'If there is at least one person who is unmoved by A Love Supreme, then Kathryn will eat her hat'. It is now clear that this is a conditional whose antecedent is a quantified expression; so we can symbolize the entire sentence with a conditional as the main logical operator: '$\exists xHx \rightarrow Ek$'.

Sentence 16 can be paraphrased, 'For all people x, if x is unmoved by A Love Supreme, then x lacks taste in music'. Or, in more natural English, it can be paraphrased by 'All bassists are rock stars'. It is best symbolized as '$\forall x(Bx \rightarrow Tx)$', just like sentence 14.

The moral is that the English words 'any' and 'anyone' should typically be symbolized 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 symbolize these sentences:

17. If everyone is moved by A Love Supreme, then Eliza is moved by A Love Supreme.
18. Everyone is such that, if she is moved by A Love Supreme, then Eliza is moved by A Love Supreme.

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

> e: Eliza

Sentence 17 is a conditional, whose antecedent is 'everyone is moved by A Love Supreme'. So we will symbolize it with '$\forall x Mx \to Me$'. This sentence is *necessarily* true: if *everyone* is indeed moved by A Love Supreme, then take any one you like—for example Eliza—and she will be moved by A Love Supreme.

Sentence 18, by contrast, might best be paraphrased by 'every person x is such that, if x is moved by A Love Supreme, then Eliza is moved by A Love Supreme'. This is symbolized by '$\forall x(Mx \to Me)$'. And this sentence is false. Kathryn is moved by A Love Supreme. So 'Mk' is true. But Eliza, tragically, is not moved by A Love Supreme, so 'Me' is false. Accordingly, '$Mk \to Me$' will be false. So '$\forall x(Mx \to Be)$' will be false as well.

In short, '$\forall x Mx \to Me$' and '$\forall x(Mx \to Me)$' 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 '$\neg Mk \to Mk$', the scope of '$\neg$' is just the antecedent of the conditional. We are saying something like: if 'Mk' is false, then 'Mk' is true. Similarly, in the sentence '$\forall x Mx \to Mk$', the scope of '$\forall x$' is just the antecedent of the conditional. We are saying something like: if 'Mx' is true of *everything*, then 'Mk' is also true.

In the sentence '$\neg(Mk \to Mk)$', the scope of '$\neg$' is the entire sentence. We are saying something like: '$(Mk \to Mk)$' is false. Similarly, in the sentence '$\forall x(Mx \to Mk)$', the scope of '$\forall x$' is the entire sentence. We are saying something like: '$(Mx \to Mk)$' 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.

symbolize each argument in FOL.

B. Using the following symbolization key:

domain: people
Kx: _____$_x$ knows the combination to the safe
Sx: _____$_x$ is a spy
Vx: _____$_x$ is a vegetarian
e: Elizabeth
p: Phillip

symbolize the following sentences in FOL:

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

C. Using this symbolization key:

domain: all animals
Ax: _____$_x$ is an alligator.
Mx: _____$_x$ is a monkey.
Rx: _____$_x$ is a reptile.
Zx: _____$_x$ lives at the zoo.
a: Amos
b: Bouncer
c: Cleo

symbolize 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 an reptile, then Amos is.
8. If any animal is an alligator, then it is a reptile.

D. For each argument, write a symbolization key and symbolize the argument in FOL.

1. Dorothy is a logician. All logicians wear funny hats. So Dorothy 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 Langston nor Ornette has been to Australia. A person could see a kangaroo only if they had been to Australia or to a zoo. Although Langston has not seen a kangaroo, Ornette has. Therefore, Ornette 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. Robin is a baby. Therefore, Robin 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 'Willie borrowed the family car from Beyoncé'. By deleting two singular terms, we can obtain any of three different two-place predicates

     Willie borrowed _____ from _____
     _____ borrowed the family car from _____
     _____ borrowed _____ from Beyoncé

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

     _____ borrowed _____ from _____

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

Now there is a little foible with the above. I have used the same symbol, '_____', to indicate a gap formed by deleting a term from a sentence. However (as Frege emphasized), these are *different* gaps. To obtain a sentence, we can

fill them in with the same term, but we can equally fill them in with different terms, and in various different orders. The following are all perfectly good sentences, and they all mean very different things:

> Langston loves Langston
> Langston loves Zora
> Zora loves Langston
> Zora loves Zora

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

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

1. Langston loves Zora.
2. Zora loves herself.
3. Langston loves Zora, but not vice versa.
4. Langston is loved by Zora.

I will start with the following representation key:

domain: people
  a: Zora
  l: Langston
 Lxy: _____$_x$ loves _____$_y$

Sentence 1 will now be symbolized by 'Lla'.

Sentence 2 can be paraphrased as 'Zora loves Zora'. It can now be symbolized by 'Laa'.

Sentence 3 is a conjunction. We might paraphrase it as 'Langston loves Zora, and Zora does not love Langston'. It can now be symbolized by '$Lla \land \neg Lal$'.

Sentence 4 might be paraphrased by 'Zora loves Langston'. It can then be symbolized by 'Lal'. Of course, this slurs over the difference in tone between the active and passive voice; such nuances are lost in FOL.

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

 Mxy: _____$_y$ loves _____$_x$

Here, we have used the same English word ('loves') as we used in our symbolization key for 'Lxy'. However, we have swapped the order of the *gaps* around (just look closely at those little subscripts!) So 'Mla' and 'Lal' now *both* symbolize 'Zora loves Langston'. 'Mal' and 'Lla' now *both* symbolize 'Langston loves Zora'. Since love can be unrequited, these are very different claims.

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

16.2 The order of quantifiers

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

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

Sentence 5 can be symbolized by '$\forall x \exists y Lxy$', and would be true of a love-triangle. For example, suppose that our domain of discourse is restricted to Zora, Louise and Langston. Suppose also that Langston loves Louise but not Zora, that Zora loves Langston but not Louise, and that Louise loves Zora but not Langston. Then sentence 5 is true.

Sentence 6 is symbolized by '$\exists y \forall x Lxy$'. Sentence 6 is *not* true in the situation just described. Again, suppose that our domain of discourse is restricted to Zora, Louise and Langston. Then this requires that all of Louise, Zora and Langston 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. | ($\forall\exists$) |
| So: | There is some truth, that no person can know. | ($\exists\forall$) |

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

| | Every cat has its day. | ($\forall\exists$) |
| So: | There is a day for all the cats. | ($\exists\forall$) |

The moral is: take great care with the order of quantification.

16.3 Stepping-stones to symbolization

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

domain: people and cats
Dx: _____$_x$ is a cat
Fxy: _____$_x$ is a friend of _____$_y$
Oxy: _____$_x$ owns _____$_y$
a: Albert

And now let's try to symbolize these sentences:

---

[1]Thanks to Rob Trueman for the example.

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

Sentence 7 can be paraphrased as, 'There is a cat that Albert owns'. This can be symbolized by '$\exists x(Dx \land Oax)$'.

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

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

I need to pause to clarify something here. In working out how to symbolize the last sentence, we wrote down '$\exists y(y$ is a cat 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 ('cat owner'). It is really is *just a stepping-stone* on the way to symbolizing 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 9 can be paraphrased as, 'Everyone who is a friend of Albert is a cat owner'. Using our stepping-stone tactic, we might write

$$\forall x \big[Fxa \to x \text{ is a cat owner}\big]$$

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

$$\forall x \big[Fxa \to \exists x(Dx \land Oxx)\big]$$

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

$$\forall x \big[Fxa \to \exists z(Dz \land Oxz)\big]$$

and this adequately symbolizes sentence 9.

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

$$\forall x \big[x \text{ is a cat owner} \to \exists y(y \text{ is a cat owner} \land Fyx)\big]$$

Completing the symbolization, we end up with

$$\forall x \big[\exists z(Dz \land Oxz) \to \exists y\big(\exists z(Dz \land Oyz) \land Fyx\big)\big]$$

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 is clear which quantifier the letter falls under. We might graphically represent the scope of the quantifiers thus:

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

   Sentence 11 is the trickiest yet. First we paraphrase it as 'For any x that is a friend of a cat owner, x owns a cat which is also owned by a friend of x'. Using our stepping-stone tactic, this becomes:

$\forall x \big[x$ is a friend of a cat owner $\to$

$\qquad\qquad\qquad\qquad x$ owns a cat which is owned by a friend of $x\big]$

Breaking this down a bit more:

$\forall x \big[\exists y(Fxy \land y$ is a cat owner$) \to$

$\qquad\qquad\qquad\qquad \exists y(Dy \land Oxy \land y$ is owned by a friend of $x)\big]$

And a bit more:

$$\forall x \big[\exists y(Fxy \land \exists z(Dz \land Oyz)) \to \exists y(Dy \land Oxy \land \exists z(Fzx \land Ozy))\big]$$

And we are done!

Practice exercises

A. Using this symbolization key:

domain: all animals
$\quad Ax$: _____$_x$ is an alligator
$\quad Mx$: _____$_x$ is a monkey
$\quad Rx$: _____$_x$ is a reptile
$\quad Zx$: _____$_x$ lives at the zoo
$\quad Lxy$: _____$_x$ loves _____$_y$
$\qquad a$: Amos
$\qquad b$: Bouncer

     c: Cleo

symbolize 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 symbolization key:

domain: all animals
    Dx: _____$_x$ is a cat
    Sx: _____$_x$ likes samurai movies
   Lxy: _____$_x$ is larger than _____$_y$
     b: Oliver
     e: Eliza
     i: Ivan

symbolize the following sentences in FOL:

1. Oliver is a cat who likes samurai movies.
2. Oliver, Eliza, and Ivan are all cats.
3. Oliver is larger than Ada, and Ivan is larger than Eliza.
4. All cats like samurai movies.
5. Only cats like samurai movies.
6. There is a cat that is larger than Eliza.
7. If there is a cat larger than Ivan, then there is a cat larger than Eliza.
8. No animal that likes samurai movies is larger than Eliza.
9. No cat is larger than Ivan.
10. Any animal that dislikes samurai movies is larger than Oliver.
11. There is an animal that is between Oliver and Eliza in size.
12. There is no cat that is between Oliver and Eliza in size.
13. No cat is larger than itself.
14. Every cat is larger than some cat.
15. There is an animal that is smaller than every cat.
16. If there is an animal that is larger than any cat, then that animal does not like samurai movies.

C. Using the following symbolization key:

domain: people and dishes at a potluck
    Rx: _____$_x$ has run out.
    Tx: _____$_x$ is on the table.
    Fx: _____$_x$ is food.

Px: _____$_x$ is a person.
Lxy: _____$_x$ likes _____$_y$.
m: Miles
n: Nina
g: the guacamole

symbolize 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 Miles does.
5. Nina only likes the dishes that have run out.
6. Nina likes no one, and no one likes Nina.
7. Miles likes anyone who likes the guacamole.
8. Miles 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 symbolization key:

domain: people
Sx: _____$_x$ sings.
Fx: _____$_x$ is female.
Mx: _____$_x$ is male.
Cxy: _____$_x$ is a child of _____$_y$.
Sxy: _____$_x$ is a sibling of _____$_y$.
m: Michael
j: Janet
e: Joe
p: Paris

symbolize the following arguments in FOL:

1. All of Joe's children are singers.
2. Janet is Joe's daughter.
3. Joe has a daughter.
4. Joe is an only child.
5. All of Joe's sons sing.
6. Janet has no daughters.
7. Paris is Janet's niece.
8. Michael is Janet's brother.
9. Joe's brothers have no children.
10. Janet is an aunt.
11. Everyone who sings has a brother who also sings.
12. Every woman who sings is the child of someone who sings.

Identity 17

Consider this sentence:

1. Nina owes money to everyone

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

Oxy: _____$_x$ owes money to _____$_y$
  n: Nina

we can symbolize sentence 1 by '$\forall x Onx$'. But this has a (perhaps) odd consequence. It requires that Nina owes money to every member of the domain (whatever the domain may be). The domain certainly includes Nina. So this entails that Nina owes money to herself.

Perhaps we meant to say:

2. Nina owes money to everyone *else*
3. Nina owes money to everyone *other than* Nina
4. Nina owes money to everyone *except* Nina herself

But we do not know how to deal with the italicized words yet. The solution is to add another symbol to FOL.

17.1 Adding identity

The symbol '$=$' is a two-place predicate. Since it is to have a special meaning, we shall write it a bit differently: we put it between two terms, rather than out front. And it *does* have a very particular meaning. We *always* adopt the following symbolization key:

$x = y$: _____$_x$ is identical to _____$_y$

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.

Now suppose we want to symbolize this sentence:

5. Nina is Eunice Waymon.

Let us add to our symbolization key:

> e: Eunice Waymon

Now sentence 5 can be symbolized as '$n = e$'. This means that the names 'n' and 'e' both name the same thing.

   We can also now deal with sentences 2–4. All of these sentences can be paraphrased as 'Everyone who is not Nina is owed money by Nina'. Paraphrasing some more, we get: 'For all x, if x is not Nina, then x is owed money by Nina'. Now that we are armed with our new identity symbol, we can symbolize this as '$\forall x(\neg x = n \rightarrow Onx)$'.

   This last sentence contains the formula '$\neg x = n$'. 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, '$x = n$'.

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

6. No one besides Nina owes money to John.
7. Only Nina owes John money.

Letting 'j' name John, sentence 6 can be paraphrased as, 'No one who is not Nina owes money to John'. This can be symbolized by '$\neg \exists x(\neg x = n \wedge Oxj)$'. Equally, sentence 6 can be paraphrased as 'for all x, if x owes money to John, then x is Nina'. Then it can be symbolized as '$\forall x(Oxj \rightarrow x = n)$'.

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

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 symbolization key:

> Ax: _____x is an apple

Sentence 8 does not require identity. It can be adequately symbolized by '$\exists x Ax$': There is some apple; perhaps many, but at least one.

   It might be tempting to also translate sentence 9 without identity. Yet consider the sentence '$\exists x \exists y(Ax \wedge Ay)$'. Roughly, this says that there is some apple x in the domain and some apple y in the domain. Since nothing precludes these from being one and the same apple, this would be true even if there were

only one apple. In order to make sure that we are dealing with *different* apples, we need an identity predicate. Sentence 9 needs to say that the two apples that exist are not identical, so it can be symbolized by '$\exists x \exists y (Ax \land Ay \land \neg x = y)$'.

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

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:

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

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 symbolized by

$$\forall x \forall y \big[(Ax \land Ay) \to x = y\big]$$

The two sentences will turn out to be logically equivalent.

In a similar way, 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:

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

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:

$$\forall x \forall y \forall z \big[(Ax \land Ay \land Az) \to (x = y \lor x = z \lor y = z)\big]$$

17.4   There are exactly...

We can now consider precise statements, 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:

$$\exists x Ax \wedge \forall x \forall y \big[(Ax \wedge Ay) \to x = y\big]$$

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:

$$\exists x \big[Ax \wedge \forall y(Ay \to x = y)\big]$$

Similarly, sentence 14 may be paraphrased as, 'There are *at least* two apples, and there are *at most* two apples'. Thus we could offer

$$\exists x \exists y(Ax \wedge Ay \wedge \neg x = y) \wedge \forall x \forall y \forall z \big[(Ax \wedge Ay \wedge Az) \to (x = y \vee x = z \vee y = z)\big]$$

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:

$$\exists x \exists y \big[Ax \wedge Ay \wedge \neg x = y \wedge \forall z(Az \to (x = z \vee y = z))\big]$$

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 symbolization key, to symbolize 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 symbolize either sentence with either of the following:

$$\exists x \exists y \neg x = y \wedge \neg \exists x \exists y \exists z(\neg x = y \wedge \neg y = z \wedge \neg x = z)$$
$$\exists x \exists y \big[\neg x = y \wedge \forall z(x = z \vee y = z)\big]$$

Practice exercises

A. Explain why:

- '$\exists x \forall y(Ay \leftrightarrow x = y)$' is a good symbolization of 'there is exactly one apple'.
- '$\exists x \exists y \big[\neg x = y \wedge \forall z(Az \leftrightarrow (x = z \vee y = z))\big]$' is a good symbolization of 'there are exactly two apples'.

Definite descriptions 18

Consider sentences like:

1. Maya Angelou is the author of *Beloved*.
2. The author of *Beloved* was from St. Louis.
3. The author of *Beloved* is the person at the podium.

These are definite descriptions: they are meant to pick out a *unique* object. They should be contrasted with *indefinite* descriptions, such as 'Maya Angelou is *an* author'. They should equally be contrasted with *generics*, such as '*The* cat is a mammal' (it's inappropriate to ask *which* cat). 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* author— whoever it is—is indeed *a* author. We want to preserve that information in our symbolization.

A second option would be to use a *new* definite description operator, such as '$\imath$'. The idea would be to symbolize 'the F' as '$\imath x F x$'; or to symbolize 'the G' as '$\imath x G x$', etc. Expression of the form $\imath x \mathcal{A} x$ would then behave like names. If we followed this path, then using the following symbolization key:

domain: people
Ax: _____$_x$ is an author of *Beloved*.
Px: _____$_x$ is at the podium.
Sx: _____$_x$ is from St Louis.
m: Maya Angelou

We could symbolize sentence 1 with '$\imath x A x = m$', sentence 2 with '$S\imath x A x$', and sentence 3 with '$\imath x A x = \imath x P x$'.

However, it would be nice if we didn't have to add a new symbol to FOL. And indeed, 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 analyzed the notion of a definite description as follows:[1]

<div align="center">

the F is G **iff** there is at least one F, *and*

there is at most one F, *and*

every F is G

</div>

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 objects in my line of sight).

If we accept Russell's analysis of definite descriptions, then we can symbolize 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:

$$\exists x Fx \land \forall x \forall y((Fx \land Fy) \to x = y) \land \forall x(Fx \to Gx)$$

In fact, we could express the same point rather more crisply, by recognizing that the first two conjuncts just amount to the claim that there is *exactly* one F, and that the last conjunct tells us that that object is F. So, equivalently, we could offer:

$$\exists x \big[Fx \land \forall y(Fy \to x = y) \land Gx \big]$$

Using these sorts of techniques, we can now symbolize sentences 1–3 without using any new-fangled fancy operator, such as '$\imath$'.

Sentence 1 is exactly like the examples we have just considered. So we would symbolize it by '$\exists x(Ax \land \forall y(Ay \to x = y) \land x = n)$'.

Sentence 2 poses no problems either: '$\exists x(Ax \land \forall y(Ay \to x = y) \land Sx)$'.

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 author, x, and there is exactly one deputy, y, and x = y'. So we can symbolize it by:

$$\exists x \exists y \big(\big[Ax \land \forall z(Az \to x = z) \big] \land \big[Py \land \forall z(Pz \to y = z) \big] \land x = y \big)$$

Note that I have made sure that the formula '$x = y$' falls within the scope of both quantifiers!

---

[1]Bertrand Russell, 'On Denoting', 1905, *Mind 14*, pp. 479–93; also Russell, *Introduction to Mathematical Philosophy*, 1919, London: Allen and Unwin, ch. 16.

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.

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),[2] suppose I, Kathryn Lindeman, claim:

4. I am dating the present king of France.

Using the following symbolization key:

k: Kathryn
Kx: _____$_x$ is a present king of France
Dxy: _____$_x$ is dating _____$_y$

Sentence 4 would be symbolized by '$\exists x(\forall y(Ky \leftrightarrow x = y) \land Dkx)$'. 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 such that he and Kathryn are dating.
6. There is a unique present King of France, but Kathryn is not dating him.

Sentence 5 might be paraphrased by 'It is not the case that: the present King of France and Kathryn are dating'. It will then be symbolized by '$\neg\exists x\big[Kx \land \forall y(Ky \to x = y) \land Dkx\big]$'. We might call this *outer* negation, since the negation governs the entire sentence. Note that it will be true if there is no present King of France.

Sentence 6 can be symbolized by '$\exists x(Kx \land \forall y(Ky \to x = y) \land \neg Dkx)$. 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 dating Kathryn.

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.

---

[2]Neale, *Descriptions*, 1990, SLU: MIT Press.

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.[3] Rather, they involve presupposition failure, and need to be regarded as *neither* true *nor* false.

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 Kathryn is dating the present King of France?[4]

Keith Donnellan raised a second sort of worry, which (very roughly) can be brought out by thinking about a case of mistaken identity.[5] Two men stand in the corner: a very tall man drinking what looks like a gin martini; and a very short man drinking what looks like a pint of water. Seeing them, Dorothy says:

7. The gin-drinker is very tall!

Russell's analysis will have us render Dorothy'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, Dorothy has said something false. But don't we want to say that Dorothy has said something *true*?

Again, one might wonder how clear our intuitions are on this case. We can all agree that Dorothy 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* Dorothy's intentions were frustrated, and so why she said something false. This is easy enough to do: Dorothy said something false because she had false beliefs about the men's drinks; if Dorothy's beliefs about the drinks had been true, then she would have said something true.[6]

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

---

[3]P.F. Strawson, 'On Referring', 1950, *Mind 59*, pp. 320–34.

[4]This is Neale's (1990) line.

[5]Keith Donnellan, 'Reference and Definite Descriptions', 1966, *Philosophical Review 77*, pp. 281–304.

[6]Interested parties should read Saul Kripke, 'Speaker Reference and Semantic Reference', 1977, in French et al (eds.), *Contemporary Perspectives in the Philosophy of Language*, Minneapolis: University of Minnesota Press, pp. 6-27.

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 symbolization key:

domain: people
Kx: _____$_x$ knows the combination to the safe.
Sx: _____$_x$ is a spy.
Vx: _____$_x$ is a vegetarian.
Txy: _____$_x$ trusts _____$_y$.
e: Elizabeth
p: Phillip

symbolize the following sentences in FOL:

1. Elizabeth trusts a vegetarian.
2. Everyone who trusts Phillip trusts a vegetarian.
3. Everyone who trusts Phillip trusts someone who trusts a vegetarian.
4. Only Phillip knows the combination to the safe.
5. Phillip trusts Elizabeth, but no one else.
6. The person who knows the combination to the safe is a vegetarian.
7. The person who knows the combination to the safe is not a spy.

B. Using the following symbolization key:

domain: cards in a standard deck
Bx: _____$_x$ is black.
Cx: _____$_x$ is a club.
Dx: _____$_x$ is a deuce.
Jx: _____$_x$ is a jack.
Mx: _____$_x$ is a man with an axe.
Ox: _____$_x$ is one-eyed.
Wx: _____$_x$ is wild.

symbolize 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 symbolization key:

domain: animals in the world
  Bx: _____$_x$ is in Farmer Brown's field.
  Hx: _____$_x$ is a horse.
  Px: _____$_x$ is a Pegasus.
  Wx: _____$_x$ has wings.

symbolize 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 symbolized 'Maya Angelou is the author of *Beloved*' by '$\exists x(Ax \land \forall y(Ay \rightarrow x = y) \land x = m)$'. Two equally good symbolizations would be:

- $Am \land \forall y(Ay \rightarrow m = y)$
- $\forall y(Ay \leftrightarrow y = m)$

Explain why these would be equally good symbolizations.

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	$A, B, C, \ldots, Z$
with subscripts, as needed	$A_1, B_1, Z_1, A_2, A_{25}, J_{375}, \ldots$
Constants	$a, b, c, \ldots, r$
with subscripts, as needed	$a_1, b_{224}, h_7, m_{32}, \ldots$
Variables	s, t, u, v, w, x, y, z
with subscripts, as needed	$x_1, y_1, z_1, x_2, \ldots$
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 will take some unpacking.

We start by defining the notion of a term.

> A TERM is any name or any variable.

So, here are some terms:

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

We next need to define atomic formulas.

> 1. If $\mathcal{R}$ is an n-place predicate and $t_1, t_2, \ldots, t_n$ are terms, then $\mathcal{R} t_1 t_2 \ldots t_n$ is an atomic formula.
> 2. If t_1 and t_2 are terms, then $t_1 = t_2$ is an atomic formula.
> 3. Nothing else is an atomic formula.

The use of swashfonts here follows the conventions laid down in §7. So, '$\mathcal{R}$' is not itself a predicate of FOL. Rather, it is a symbol of our metalanguage (augmented English) that we use to talk about any predicate of FOL. Similarly, 't_1' is not a term of FOL, but a symbol of the metalanguage that we can use to talk about any term of FOL. So, where 'F' is a one-place predicate, 'G' is a three-place predicate, and 'S' is a six-place predicate, here are some atomic formulas:

$$x = a$$
$$a = b$$
$$Fx$$
$$Fa$$
$$Gxay$$
$$Gaaa$$
$$Sx_1 x_2 abyx_1$$
$$Sby_{254}zaaz$$

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 $\mathcal{A}$ is a formula, then $\neg\mathcal{A}$ is a formula.

3. If $\mathcal{A}$ and $\mathcal{B}$ are formulas, then $(\mathcal{A} \wedge \mathcal{B})$ is a formula.

4. If $\mathcal{A}$ and $\mathcal{B}$ are formulas, then $(\mathcal{A} \vee \mathcal{B})$ is a formula.

5. If $\mathcal{A}$ and $\mathcal{B}$ are formulas, then $(\mathcal{A} \rightarrow \mathcal{B})$ is a formula.

6. If $\mathcal{A}$ and $\mathcal{B}$ are formulas, then $(\mathcal{A} \leftrightarrow \mathcal{B})$ is a formula.

7. If $\mathcal{A}$ is a formula, χ is a variable, $\mathcal{A}$ contains at least one occurrence of χ, and $\mathcal{A}$ contains neither $\forall\chi$ nor $\exists\chi$, then $\forall\chi\mathcal{A}$ is a formula.

8. If $\mathcal{A}$ is a formula, χ is a variable, $\mathcal{A}$ contains at least one occurrence of χ, and $\mathcal{A}$ contains neither $\forall\chi$ nor $\exists\chi$, then $\exists\chi\mathcal{A}$ is a formula.

9. Nothing else is a formula.

So, assuming again that 'F' is a one-place predicate, 'G' is a three-place predicate and 'H' is a six place-predicate, here are some formulas:

$$Fx$$
$$Gayz$$
$$Syzyayx$$
$$(Gayz \rightarrow Syzyayx)$$
$$\forall z(Gayz \rightarrow Syzyayx)$$
$$Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx)$$
$$\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx))$$
$$\forall x\exists y(Fx \leftrightarrow \forall z(Gayz \rightarrow Syzyayx))$$

But this is *not* a formula:

$$\forall x\exists x Gxxx$$

Certainly '$Gxxx$' is a formula. And '$\exists x Gxxx$' 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: '$\exists x Gxxx$' 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 last, 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:

$$\overbrace{\forall x\, \overbrace{\exists y(Fx \leftrightarrow \overbrace{\forall z(Gayz \rightarrow Syzyayx)}^{\text{scope of `}\forall z\text{'}}}^{\text{scope of `}\exists y\text{'}}}^{\text{scope of `}\forall x\text{'}})$$

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 symbolization key:

domain: people
  Lxy: _____x loves _____y
     b: Beyoncé

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 Beyoncé (the person named by 'b') loves herself. *But 'z' is a variable, and does not name anyone or any thing.*

   Of course, if we put an existential quantifier out front, obtaining '$\exists z Lzz$', then this would be true iff someone loves herself. Equally, if we wrote '$\forall z Lzz$', this would be true iff everyone loves herself. The point is that we need a quantifier to tell us how to deal with a variable.

   Let's make this idea precise.

> A BOUND VARIABLE is an occurrence of a variable χ that is within the scope of either $\forall\chi$ or $\exists\chi$.
>
> A FREE VARIABLE is any variable that is not bound.

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 '$(Ex \to Lzx)$', 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.

Practice exercises

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

1. $\exists x Lxy \land \forall y Lyx$
2. $\forall x Ax \land Bx$
3. $\forall x(Ax \land Bx) \land \forall y(Cx \land Dy)$
4. $\forall x \exists y[Rxy \to (Jz \land Kx)] \lor Ryx$
5. $\forall x_1(Mx_2 \leftrightarrow Lx_2x_1) \land \exists x_2 Lx_3x_2$

Chapter 5

interpretations

Extensionality 20

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 symbolization key, e.g.:

> P: The Arch is in St. Louis

But recall from §9 that this should be taken to mean:

- The TFL sentence 'P' is to take the same truth value as the English sentence 'The Arch is in St. Louis' (whatever that truth value may be)

The point that I emphasized is that TFL cannot handle differences in meaning that go beyond mere differences in truth value.

20.1   symbolizing 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 quantifier expressions. This enables us to consider what is *true of* some particular object, or of some or all objects. But we can do no more than that.

When we provide a symbolization key for some FOL predicates, such as:

> Cx: _____$_x$ teaches PHIL405 at SLU in the Fall of 2017

we do not carry the *meaning* of the English predicate across into our FOL predicate. We are simply stipulating something like the following:

- 'Cx' and '_____$_x$ teaches PHIL405 at SLU in the Fall of 2017' are to be *true of* exactly the same things.

So, in particular:

- 'Cx' is to be true of all and only those things which teach PHIL405 at SLU in the Fall of 2017 (whatever those things might be).

This is an indirect stipulation. Alternatively we can stipulate predicate extensions directly. We can stipulate that 'Cx' is to be true of Kathryn Lindeman, and Kathryn Lindeman alone. As it happens, this direct stipulation would have the same effect as the indirect stipulation. But note that the English predicates '_____ is Kathryn Lindeman' and '_____ teaches PHIL405 at SLU in the Fall of 2017' have very different meanings!

The point is that FOL does not give us any 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.

For this reason, I say only that FOL sentences *symbolize* English sentences. It is doubtful that we are *translating* English into FOL, for translations should preserve meanings.

20.2 A word on extensions

We can stipulate directly what predicates are to be true of. So it is worth noting that our stipulations can be as arbitrary as we like. For example, we could stipulate that 'Hx' should be true of, and only of, the following objects:

<div align="center">
Etta James

the number π

every top-F key on every piano ever made
</div>

Now, the objects that we have listed have nothing particularly in common. But this doesn't matter. Logic doesn't care about what strikes us mere humans as 'natural' or 'similar'. Armed with this interpretation of 'Hx', suppose I now add to my symbolization key:

  e: Etta James
  m: Mars
  p: the number π

Then 'He' and 'Hp' will both be true, on this interpretation, but 'Hm' will be false, since Mars was not among the stipulated objects.

(This process of explicit stipulation is sometimes described as stipulating the *extension* of a predicate.)

20.3 Many-place predicates

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

  Lxy: _____$_x$ loves _____$_y$

Given what I said above, this symbolization key should be read as saying:

- 'Lxy' and '_____$_x$ loves _____$_y$' are to be true of exactly the same things

So, in particular:

- 'Lxy' is to be true of x and y (in that order) iff x loves y.

It is important that we insist upon the order here, since love—famously—is not always reciprocated. (Note that 'x' and 'y' here are symbols of augmented English, and that they are being *used*. By contrast, 'x' and 'y' are symbols of FOL, and they are being *mentioned*.)

That is an indirect stipulation. What about a direct stipulation? This is slightly harder. If we *simply* list objects that fall under 'Lxy', we will not know whether they are the lover or the beloved (or both). We have to find a way to include the order in our explicit stipulation.

To do this, we can specify that two-place predicates are true of *pairs* of objects, where the order of the pair is important. Thus we might stipulate that 'Bxy' is to be true of, and only of, the following pairs of objects:

<Langton, Maya>
<Etta, Nina>
<Nina, Etta>

Here the angle-brackets keep us informed concerning order. Suppose I now add the following stipulations:

l: Langton
m: Maya
h: Etta
s: Nina

Then 'Blm' will be true, since <Langton, Maya> was in my explicit list. But 'Bml' will be false, since <Maya, Langton> was not in my list. However, both 'Bhs' and 'Bsh' will be true, since both <Etta, Nina> and <Nina, Etta> are in my explicit list

To make these ideas more precise, we would need to develop some *set theory*. This will give you some apparatus for dealing with extensions and with ordered pairs (and ordered triples, etc.) However, set theory is not covered in this book. So I will 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:

$$Aa \leftrightarrow Ab$$
$$Ba \leftrightarrow Bb$$
$$Raa \leftrightarrow Rbb$$
$$Raa \leftrightarrow Rab$$
$$Rca \leftrightarrow Rcb$$
$$\forall x Rxa \leftrightarrow \forall x Rxb$$

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 two objects, then they are really just one and the same object after all. 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: Janet Jackson, Nina Simone
- a: Janet Jackson
- b: Nina Simone
- • 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, Janet Jackson and Nina Simone are two distinct people, not one and the same!
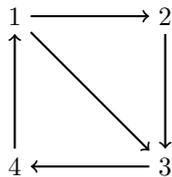
20.5   Interpretation

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—other than '='—a specification of what things (in what order) the predicate is to be true of

The symbolization keys that I considered in chapter 4 consequently give us one very convenient way to present an interpretation. We will continue to use them throughout this chapter. However, it is sometimes also convenient to present an interpretation *diagrammatically*.

Suppose we want to consider just a single two-place predicate, 'Rxy'. Then we can represent it just by drawing an arrow between two objects, and stipulate that 'Rxy' is to hold of x and y just in case there is an arrow running from x to y in our diagram. As an example, we might offer:

This would be suitable to characterize an interpretation whose domain is the first four positive whole numbers, and which interprets 'Rxy' as being true of and only of:

$$<1, 2>, <2, 3>, <3, 4>, <4, 1>, <1, 3>$$

Equally we might offer:
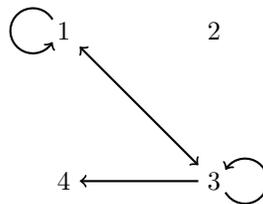
for an interpretation with the same domain, which interprets 'Rxy' as being true of and only of:

$$<1, 3>, <3, 1>, <3, 4>, <1, 1>, <3, 3>$$

If we wanted, we could make our diagrams more complex. For example, we could add names as labels for particular objects. Equally, to symbolize the extension of a one-place predicate, we might simply draw a ring around some particular objects and stipulate that the thus encircled objects (and only them) are to fall under the predicate 'Hx', say.

Truth in FOL   21

We know what interpretations are. Since, among other things, they tell us which predicates are true of which objects, they will provide us with an account of the truth of atomic sentences. But we must also present a detailed account of what it is for an arbitrary 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

We need to explain truth for all three kinds of sentence.

I will offer a completely general explanation in this section. However, to try to keep the explanation comprehensible, I will at several points use the following interpretation:

domain: all people born before 2000CE
a: Aristotle
b: Beyoncé
Sx: _____x is a singer
Rxy: _____x was born before _____y

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 'Sb' should be true just in case 'Sx' is true of 'b'. Given our go-to interpretation, this is true iff Beyoncé is singer. Beyoncé is a singer. So the sentence is true. Equally, 'Sa' is false on our go-to interpretation.

Likewise, on this interpretation, 'Rba' is true iff the object named by 'b' was born before the object named by 'a'. Well, Beyoncé was not born before Aristotle. So 'Rba' is false. Obviously, 'Rbb' is true: Beyoncé was not born before Beyoncé (which follows from the more general fact that no one can be born before they are born).

Dealing with atomic sentences, then, is very intuitive. When $\mathcal{R}$ is an n-place predicate and $a_1, a_2, \ldots, a_n$ are names,

> $\mathcal{R}a_1 a_2 \ldots a_n$ is true in an interpretation **iff**
> $\mathcal{R}$ is true of the objects named by $a_1, a_2, \ldots, a_n$ in that interpretation (considered in that order)

Recall, though, that there is a second kind of atomic sentence, recall: two names connected by an identity sign constitute an atomic sentence. This kind of atomic sentence is also easy to handle. Where a and b are any names,

> $a = b$ is true in an interpretation **iff**
> a and b name the very same object in that interpretation

So in our go-to interpretation, '$b = a$' is false, since Beyoncé is distinct from Aristotle.

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 truth-functional connectives are *exactly* the same as they were when we considered TFL. Here they are:

> $\mathcal{A} \land \mathcal{B}$ is true in an interpretation **iff**
> both $\mathcal{A}$ is true and $\mathcal{B}$ is true in that interpretation
>
> $\mathcal{A} \lor \mathcal{B}$ is true in an interpretation **iff**
> either $\mathcal{A}$ is true or $\mathcal{B}$ is true in that interpretation
>
> $\neg\mathcal{A}$ is true in an interpretation **iff**
> $\mathcal{A}$ is false in that interpretation
>
> $\mathcal{A} \to \mathcal{B}$ is true in an interpretation **iff**
> either $\mathcal{A}$ is false or $\mathcal{B}$ is true in that interpretation
>
> $\mathcal{A} \leftrightarrow \mathcal{B}$ is true in an interpretation **iff**
> $\mathcal{A}$ has the same truth value as $\mathcal{B}$ in that interpretation

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

- '$a = a \land Sb$' is true
- '$Rab \land Sa$' is false because, although 'Rab' is true, 'Sa' is false
- '$a = b \lor Sb$' is true
- '$\neg a = b$' is true
- '$Sb \land \neg(a = b \land Rab)$' is true, because '$Sb$' is true and '$a = b$' is false

Make sure you understand these examples.

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 '$\forall x Fx$' 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 '$\forall x \exists y Lxy$' is true just in case '$\exists y Lxy$' is true of everything in the domain. And this is problematic, since our interpretation does not directly specify what '$\exists y Lxy$' is to be true of. Instead, whether or not this is true of something should follow just from the interpretation of 'Lxy', the domain, and the meanings of the quantifiers.

So here is a naïve second thought. We might try to say that '$\forall x \exists y Lxy$' is to be true in an interpretation iff $\exists y L\boldsymbol{a}y$ is true for *every* name $\boldsymbol{a}$ that we have included in our interpretation. And similarly, we might try to say that $\exists y L\boldsymbol{a}y$ is true just in case $L\boldsymbol{a}\boldsymbol{b}$ is true for *some* name $\boldsymbol{b}$ that we have included in our interpretation.

Unfortunately, this is not right either. To see this, observe that in our go-to interpretation, we have only given interpretations for *two* names, 'a' and 'b'. But the domain—all people born before the year 2000ce—contains many more than two people. I have no intention of trying to name *all* of them!

So here is a third thought. (And this thought is not naïve, but correct.) Although it is not the case that we have named *everyone*, each person *could* have been given a name. So we should focus on this possibility of extending an interpretation, by adding a new name. I will offer a few examples of how this might work, centering on our go-to interpretation, and I will then present the formal definition.

In our go-to interpretation, '$\exists x Rbx$' should be true. After all, in the domain, there is certainly someone who was born after Beyoncé. Rhianna is one of those people. Indeed, if we were to extend our go-to interpretation by adding the name 'h' to refer to Rhianna, then 'Rbh' would be true on this extended interpretation. And this, surely, should suffice to make '$\exists x Rbx$' true on the original go-to interpretation.

In our go-to interpretation, '$\exists x (Sx \wedge Rxb)$' should also be true. After all, in the domain, there is certainly someone who was both a singer and born before Beyoncé. Nina Simone is one such person. Indeed, if we were to extend our go-to interpretation by letting a new name, 'n', denote Nina Simone, then '$Sn \wedge Rnb$' would be true on this extended interpretation. Again, this should surely suffice to make '$\exists x (Sx \wedge Rxb)$' true on the original go-to interpretation.

In our go-to interpretation, '$\forall x \exists y Rxy$' 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, 'c', denote that person, then we would not be able to find anyone else in the domain to denote with some further new name, perhaps 'd', in such a way that 'Rcd' would

be true. Indeed, no matter *whom* we named with '*d*', '*Rcd*' would be false. And this observation is surely sufficient to make '∃*yRcy*' *false* in our extended interpretation. And this is surely sufficient to make '∀*x*∃*yRxy*' false on the original go-to interpretation.

If you have understood these three examples, then that's what matters. Strictly speaking, though, we still need to give a precise definition of the truth conditions for quantified sentences. The result, sadly, is a bit ugly, and requires a few new definitions. Brace yourself!

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

$$\mathcal{A}(\ldots \chi \ldots \chi \ldots)$$

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

$$\mathcal{A}(\ldots c \ldots c \ldots)$$

for the formula obtained by replacing every occurrence of χ in $\mathcal{A}$ with c. The resulting formula is called a SUBSTITUTION INSTANCE of $\forall \chi \mathcal{A}$ and $\exists \chi \mathcal{A}$. 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 $\forall \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$ will be true iff $\mathcal{A}(\ldots c \ldots c \ldots)$ is true no matter what object (in the domain) we name with c. Similarly, the sentence $\exists \chi \mathcal{A}$ will be true iff there is *some* way to assign the name c to an object that makes $\mathcal{A}(\ldots c \ldots c \ldots)$ true. More precisely, we stipulate:

> $\forall \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$ is true in an interpretation **iff**
> $\mathcal{A}(\ldots c \ldots c \ldots)$ is true in *every* interpretation that extends the original interpretation by assigning an object to any name c (without changing the interpretation in any other way).
>
> $\exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$ is true in an interpretation **iff**
> $\mathcal{A}(\ldots c \ldots c \ldots)$ is true in *some* interpretation that extends the original interpretation by assigning an object to some name c (without changing the interpretation in any other way).

To be clear: all this is doing is formalizing (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 Etta and Nina
- 'Ax' is to be true of both Etta and Nina
- 'Bx' is to be true of Nina only
- 'Nx' is to be true of no one
- 'w' is to refer to Etta

Determine whether each of the following sentences is true or false in that interpretation:

1. Bw
2. $Ac \leftrightarrow \neg Nw$
3. $Nc \rightarrow (Ac \vee Bc)$
4. $\forall x Ax$
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. $\exists x Bx \rightarrow \forall x Ax$
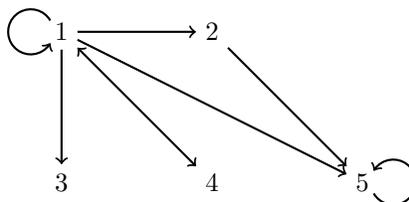
B. Consider the following interpretation:

- The domain comprises only Ornette, Miles and John
- 'Gx' is to be true of Ornette, Miles and John.
- 'Hx' is to be true of and only of Miles
- 'Mx' is to be true of and only of Ornette and John
- 'm' is to refer to Miles
- 'j' is to refer to John

Determine whether each of the following sentences is true or false in that interpretation:

1. Hm
2. Hj
3. $Mm \vee Mj$
4. $Gm \vee \neg Gm$
5. $Mm \rightarrow Gm$
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. $\exists x Gx \land \exists x \neg Gx$
15. $\forall x \exists y (Gx \land Hy)$

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. $\forall x \forall y \forall z ((Rxy \land Ryz) \to Rxz)$
6. $\forall x \forall y \forall z ((Rxy \land Rxz) \to Ryz)$
7. $\exists x \forall y \neg Rxy$
8. $\forall x (\exists y Rxy \to \exists y Ryx)$
9. $\exists x \exists y (\neg x = y \land Rxy \land Ryx)$
10. $\exists x \forall y (Rxy \leftrightarrow x = y)$
11. $\exists x \forall y (Ryx \leftrightarrow x = y)$
12. $\exists x \exists y (\neg x = y \land Rxy \land \forall z (Rzx \leftrightarrow y = z))$

Semantic concepts 22

Offering a precise definition of truth in FOL was more than a little fiddly. But now that we are done, we can define various central logical notions. These will look very similar to the definitions we offered for TFL. However, remember that they concern *interpretations*, rather than valuations.

We will use the symbol '⊨' for FOL much as we did for TFL. So:

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

means that there is no interpretation in which all of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ are true and in which $\mathcal{C}$ is false. Derivatively,

$$\vDash \mathcal{A}$$

means that $\mathcal{A}$ is true in every interpretation.

An FOL sentence $\mathcal{A}$ is a LOGICAL TRUTH iff $\mathcal{A}$ is true in every interpretation; i.e., $\vDash \mathcal{A}$.

$\mathcal{A}$ is a CONTRADICTION iff $\mathcal{A}$ is false in every interpretation; i.e., $\vDash \neg\mathcal{A}$.

$\mathcal{A}_1, \mathcal{A}_2, \ldots \mathcal{A}_n \therefore \mathcal{C}$ is VALID IN FOL iff there is no interpretation in which all of the premises are true and the conclusion is false; i.e., $\mathcal{A}_1, \mathcal{A}_2, \ldots \mathcal{A}_n \vDash \mathcal{C}$. It is INVALID IN FOL otherwise.

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

The FOL sentences $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ are JOINTLY CONSISTENT iff there is some interpretation in which all of the sentences are true. They are JOINTLY INCONSISTENT iff there is no such interpretation.

Using interpretations 23

23.1 Logical truths and contradictions

Suppose we want to show that '$\exists x Axx \to Bd$' 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 '$\exists x Axx \to Bd$' to be false, the antecedent ('$\exists x Axx$') 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 will start with a domain that has just one member. For concreteness, let's say it is the city of St. Louis.

domain: St. Louis

The name 'd' must name something in the domain, so we have no option but:

> d: St. Louis

Recall that we want '$\exists x Axx$' 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:

> Axy: _____$_x$ is identical with _____$_y$

Now 'Add' is true, so it is surely true that '$\exists x Axx$'. Next, we want 'Bd' to be false, so the referent of 'd' must not be in the extension of 'B'. We might simply offer:

> Bx: _____$_x$ is in Germany

Now we have an interpretation where '$\exists x Axx$' is true, but where 'Bd' is false. So there is an interpretation where '$\exists x Axx \to Bd$' is false. So '$\exists x Axx \to Bd$' is not a logical truth.

We can just as easily show that '$\exists x Axx \to Bd$' is not a contradiction. We need only specify an interpretation in which '$\exists x Axx \to Bd$' is true; i.e., an interpretation in which either '$\exists x Axx$' is false or 'Bd' is true. Here is one:

domain: St. Louis
     d: St. Louis
  Axy: _____$_x$ is identical with _____$_y$
   Bx: _____$_x$ is in the United States

This shows that there is an interpretation where '$\exists x Axx \to Bd$' is true. So '$\exists x Axx \to Bd$' is not a contradiction.

23.2 Logical equivalence

Suppose we want to show that '$\forall x Sx$' and '$\exists x Sx$' 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 will need at least two objects. (If we chose a domain with only one member, the two sentences would end up with the same truth value. In order to see why, try constructing some partial interpretations with one-member domains.) For concreteness, let's take:

domain: Ornette Coleman, Miles Davis

We can make '$\exists x Sx$' true by including something in the extension of 'S', and we can make '$\forall x Sx$' false by leaving something out of the extension of 'S'. For concreteness we will offer:

   Sx: _____$_x$ plays saxophone

Now '$\exists x Sx$' is true, because 'Sx' is true of Ornette Coleman. Slightly more precisely, extend our interpretation by allowing 'c' to name Ornette Coleman. 'Sc' is true in this extended interpretation, so '$\exists x Sx$' was true in the original interpretation. Similarly, '$\forall x Sx$' is false, because 'Sx' 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 '$\forall x Sx$' was false in the original interpretation. We have provided a counter-interpretation to the claim that '$\forall x Sx$' and '$\exists x Sx$' are logically equivalent.

> To show that $\mathcal{A}$ is not a logical truth, it suffices to find an interpretation where $\mathcal{A}$ is false.
> To show that $\mathcal{A}$ is not a contradiction, it suffices to find an interpretation where $\mathcal{A}$ is true.
> To show that $\mathcal{A}$ and $\mathcal{B}$ are not logically equivalent, it suffices to find an interpretation where one is true and the other is false.

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, Joseph McCarthy
   Gx: _____$_x$ hated communism
     a: Karl Marx

Given that Marx wrote *The Communist Manifesto*, 'Ga' is plainly false in this interpretation. But Joseph McCarthy was a US Senator known for attacking his opponents as communist sympathizers. So '$\exists xGx$' is true in this interpretation. Hence '$\exists xGx \rightarrow Ga$' is false, as required.

But does this interpretation make the premise true? Yes it does! Note that '$Ga \rightarrow Ga$' is true. (Indeed, it is a logical truth.) But then certainly '$\exists x(Gx \rightarrow Ga)$' is true. So the premise is true, and the conclusion is false, in this interpretation. The argument is therefore invalid.

In passing, note that we have also shown that '$\exists x(Gx \rightarrow Ga)$' does *not* entail '$\exists xGx \rightarrow Ga$'. And equally, we have shown that the sentences '$\exists x(Gx \rightarrow Ga)$' and '$\neg(\exists xGx \rightarrow Ga)$' are jointly consistent.

Let's consider a second example. Consider:

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

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

domain: US citizens currently married to another US citizen
   Lxy: _____$_x$ is married to _____$_y$

The premise is clearly true on this interpretation. Anyone in the domain is a US citizen married to some other US 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 to whom they are married. Hence '$\forall x\exists yLxy$' is true. But the conclusion is clearly false, for that would require that there is some single person who married to 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 will mix things up a bit. In §20, I described how we can present some interpretations using diagrams. For example:
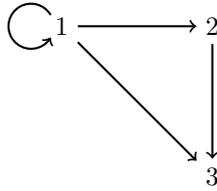
Using the conventions employed in §20, the domain of this interpretation is the first three positive whole numbers, and 'Rxy' is true of x and y just in case there is an arrow from x to y in our diagram. Here are some sentences that the interpretation makes true:

- '$\forall x \exists y Ryx$'
- '$\exists x \forall y Rxy$'   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
- '$\exists x (\exists y Ryx \wedge \neg \exists y Rxy)$'   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.:

$$\forall x \exists y Ryx, \exists x \forall y Rxy \therefore \forall x \exists y Rxy$$

$$\exists x \forall y Rxy, \exists x \forall y \neg Rxy \therefore \neg \exists x \exists y \exists z (\neg y = z \wedge Rxy \wedge Rzx)$$

and many more besides.

To show that $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \therefore \mathcal{C}$ is invalid, it suffices to find an interpretation where all of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ are true and where $\mathcal{C}$ is false.

That same interpretation will show that $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ do not entail $\mathcal{C}$.

That same interpretation will show that $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n, \neg \mathcal{C}$ are jointly consistent.

When you provide an interpretation to refute a claim—to logical truth, say, or to entailment—this is sometimes called providing a *counter-interpretation* (or providing a *counter-model*).

Practice exercises

A. Show that each of the following is neither a logical truth nor a contradiction:

1. $Da \wedge Db$
2. $\exists x Txh$
3. $Pm \wedge \neg \forall x Px$
4. $\forall z Jz \leftrightarrow \exists y Jy$

5. $\forall x(Sxmn \lor \exists yLxy)$
6. $\exists x(Gx \to \forall yMy)$
7. $\exists x(x = h \land x = i)$

B. Show that the following pairs of sentences are not logically equivalent.

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

C. Show that the following sentences are jointly consistent:

1. $Ma, \neg Na, Pa, \neg Qa$
2. $Lee, Leg, \neg Lge, \neg Lgg$
3. $\neg(Ma \land \exists xAx), Ma \lor Fa, \forall x(Fx \to Ax)$
4. $Ma \lor Mb, Ma \to \forall x\neg Mx$
5. $\forall yGy, \forall x(Gx \to Hx), \exists y\neg Iy$
6. $\exists x(Bx \lor Ax), \forall x\neg Cx, \forall x[(Ax \land Bx) \to Cx]$
7. $\exists xXx, \exists xYx, \forall x(Xx \leftrightarrow \neg Yx)$
8. $\forall x(Px \lor Qx), \exists x\neg(Qx \land Px)$
9. $\exists z(Nz \land Ozz), \forall x\forall y(Oxy \to Oyx)$
10. $\neg\exists x\forall yRxy, \forall x\exists yRxy$
11. $\neg Raa, \forall x(x = a \lor Rxa)$
12. $\forall x\forall y\forall z(x = y \lor y = z \lor x = z), \exists x\exists y \neg x = y$
13. $\exists x\exists y(Zx \land Zy \land x = y), \neg Zd, d = e$

D. Show that the following arguments are invalid:

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

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.)

Here is another bit of reasoning. Consider the sentence '$\forall x(Rxx \leftrightarrow Rxx)$'. Again, it should obviously be a logical truth. But to say precisely why is quite a challenge. We cannot say that '$Rxx \leftrightarrow Rxx$' is true in every interpretation, since '$Rxx \leftrightarrow Rxx$' is not even a *sentence* of FOL (remember that 'x' is a variable, not a name). So we have to be a bit cleverer.

> Consider some arbitrary interpretation. Consider some arbitrary member of the model's domain, which, for convenience, we will call *obbie*, and suppose we extend our original interpretation by adding a new name, 'w', to name *obbie*. Then either 'Rcw' will be true or it will be false. If 'Rcw' is true, then '$Rcc \leftrightarrow Rcw$' is true. If 'Rcw' is false, then '$Rcc \leftrightarrow Rcw$' will be true. So either way, '$Rcc \leftrightarrow Rcw$' 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 'w' to name some new object, '$Rcc \leftrightarrow Rcw$' 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! Let's take just one more example. Here is an argument which is obviously valid:

$$\forall x(Hx \land Jx) \therefore \forall x Hx$$

After all, if everything is both H and J, then everything is H. But we can only show that the argument is valid by considering what must be true in every interpretation in which the premise is true. And to show this, we would have to reason as follows:

> Consider an arbitrary interpretation in which the premise '$\forall x(Hx \land Jx)$' is true. It follows that, however we expand the interpretation

> with a new name, for example 'w', '$Hc \land Jw$' will be true in this new interpretation. 'Hw' will, then, also be true in this new interpretation. But since this held for *any* way of expanding the interpretation, it must be that '$\forall x Hx$' is true in the old interpretation. And we assumed nothing about the interpretation except that it was one in which '$\forall x(Hx \land Jx)$' is true. So any interpretation in which '$\forall x(Hx \land Jx)$' is true is one in which '$\forall x Hx$' is true. The argument is valid!

Even for a simple argument like this one, the reasoning is somewhat complicated. For longer arguments, the reasoning can be extremely torturous.

The following table summarizes 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	consider all interpretations
valid?	all interpretations	one counter-interpretation
entailment?	all interpretations	one counter-interpretation

This might usefully be compared with the table at the end of §13. The key difference resides in the fact that 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.

Chapter 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 do not necessarily give us much *insight*. Consider two arguments in TFL:

$$P \lor Q, \neg P \therefore Q$$
$$P \to Q, P \therefore Q$$

Clearly, these are valid arguments. You can confirm that they are valid by constructing four-line truth tables. But we might say that 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:

$$A_1 \to C_1 \therefore (A_1 \land A_2 \land A_3 \land A_4 \land A_5) \to (C_1 \lor C_2 \lor C_3 \lor C_4 \lor C_5)$$

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:

$$A_1 \to C_1 \therefore (A_1 \land A_2 \land A_3 \land A_4 \land A_5 \land A_6 \land A_7 \land A_8 \land A_9 \land A_{10}) \to$$
$$(C_1 \lor C_2 \lor C_3 \lor C_4 \lor C_5 \lor C_6 \lor C_7 \lor C_8 \lor C_9 \lor C_{10})$$

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 will not follow this path. We will, 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 will try to select a few basic rules of inference. Some of these will govern the behavior of the sentential connectives. Others will govern the behavior 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 will consider is based largely around work by Frederic Fitch (first published in 1952).

Basic rules for TFL

We will develop a NATURAL DEDUCTION system. For each connective, there will be INTRODUCTION rules, that allow us to prove a sentence that has that connective as the main 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 will simply call these 'proofs', but you should be aware that there are *informal proofs* too.)

As an illustration, consider:

$$\neg(A \lor B) \therefore \neg A \land \neg B$$

We will start a proof by writing the premise:

$$1 \quad \big| \; \neg(A \lor B)$$

Note that we have numbered the premise, since we will 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 '$\neg A \land \neg B$'; so we are hoping ultimately to conclude our proof with

$$n \quad \big| \; \neg A \land \neg B$$

for some number n. It doesn't matter, which line we end on, but we would obviously prefer a short proof to a long one.

Similarly, suppose we wanted to consider:

$$A \lor B, \neg(A \land C), \neg(B \land \neg D) \therefore \neg C \lor D$$

The argument has three premises, so we start by writing them all down, numbered, and drawing a line under them:

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

and we are hoping to conclude with some line:

$$
\begin{array}{c|l}
n & \neg C \vee D \\
\end{array}
$$

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 Willie is both stoner and a Texan. One obvious way to do this would be as follows: first I show that Willie is stoner; then I show that Willie is a Texan; 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 symbolization key:

S: Willie is a stoner
T: Willie is a Texan

Perhaps I am working through a proof, and I have obtained 'S' on line 8 and 'T' on line 15. Then on any subsequent line I can obtain '$S \wedge T$' thus:

$$
\begin{array}{c|l}
8 & S \\
15 & T \\
 & S \wedge T \qquad \wedge\text{I } 8,\ 15 \\
\end{array}
$$

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:

$$
\begin{array}{c|l}
8 & S \\
15 & T \\
 & T \wedge S \qquad \wedge\text{I } 15,\ 8 \\
\end{array}
$$

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

$$
\begin{array}{c|l}
m & \mathcal{A} \\[4pt]
n & \mathcal{B} \\[4pt]
  & \mathcal{A} \land \mathcal{B} \qquad \land\mathrm{I}\ m,\,n
\end{array}
$$

To be clear, the statement of the rule is *schematic*. It is not itself a proof. '$\mathcal{A}$' and '$\mathcal{B}$' are not sentences of 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 emphasize 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 '$\land$' into our proof where it may have been absent. Correspondingly, we have a rule that *eliminates* that symbol. Suppose you have shown that Willie is both a stoner and a Texan. You are entitled to conclude that Willie is a stoner. Equally, you are entitled to conclude that Willie is a Texan. Putting this together, we obtain our conjunction elimination rule(s):

$$
\begin{array}{c|l}
m & \mathcal{A} \land \mathcal{B} \\[4pt]
  & \mathcal{A} \qquad \land\mathrm{E}\ m
\end{array}
$$

and equally:

$$
\begin{array}{c|l}
m & \mathcal{A} \land \mathcal{B} \\[4pt]
  & \mathcal{B} \qquad \land\mathrm{E}\ m
\end{array}
$$

The point is simply that, when you have a conjunction on some line of a proof, you can obtain either of the conjuncts by $\land$E. (One point, might be worth emphasizing: you can only apply this rule when conjunction is the main logical operator. So you cannot infer 'D' just from '$C \lor (D \land E)$'!)

Even with just these two rules, we can start to see some of the power of our formal proof system. Consider:

$$
[(A \lor B) \to (C \lor D)] \land [(E \lor F) \to (G \lor H)]
$$
$$
\therefore\ [(E \lor F) \to (G \lor H)] \land [(A \lor B) \to (C \lor D)]
$$

The main logical operator in both the premise and conclusion of this argument is '∧'. 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:

$$1 \quad [(A \lor B) \to (C \lor D)] \land [(E \lor F) \to (G \lor H)]$$

From the premise, we can get each of the conjuncts by ∧E. The proof now looks like this:

$$
\begin{array}{lll}
1 & [(A \lor B) \to (C \lor D)] \land [(E \lor F) \to (G \lor H)] & \\
2 & [(A \lor B) \to (C \lor D)] & \land\text{E } 1 \\
3 & [(E \lor F) \to (G \lor H)] & \land\text{E } 1
\end{array}
$$

So by applying the ∧I rule to lines 3 and 2 (in that order), we arrive at the desired conclusion. The finished proof looks like this:

$$
\begin{array}{lll}
1 & [(A \lor B) \to (C \lor D)] \land [(E \lor F) \to (G \lor H)] & \\
2 & [(A \lor B) \to (C \lor D)] & \land\text{E } 1 \\
3 & [(E \lor F) \to (G \lor H)] & \land\text{E } 1 \\
4 & [(E \lor F) \to (G \lor H)] \land [(A \lor B) \to (C \lor D)] & \land\text{I } 3, 2
\end{array}
$$

This is a very simple proof, but it shows how we can chain rules of proof together into longer proofs. In passing, note that investigating this argument with a truth table would have required a staggering 256 lines; our formal proof required only four lines.

It is worth giving another example. Way back in §10.3, we noted that this argument is valid:

$$A \land (B \land C) \therefore (A \land B) \land C$$

To provide a proof corresponding with this argument, we start by writing:

$$1 \quad A \land (B \land C)$$

From the premise, we can get each of the conjuncts by applying ∧E twice. And we can then apply ∧E twice more, so our proof looks like:

$$
\begin{array}{lll}
1 & A \land (B \land C) & \\
2 & A & \land\text{E } 1 \\
3 & B \land C & \land\text{E } 1 \\
4 & B & \land\text{E } 3 \\
5 & C & \land\text{E } 3
\end{array}
$$

But now we can merrily reintroduce conjunctions in the order we wanted them, so that our final proof is:

1	$A \land (B \land C)$	
2	A	$\land$E 1
3	$B \land C$	$\land$E 1
4	B	$\land$E 3
5	C	$\land$E 3
6	$A \land B$	$\land$I 2, 4
7	$(A \land B) \land C$	$\land$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 will not do this. Instead, we will 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 $\land$I rule, there is no need to apply it to different sentences. So we can formally prove 'A' from 'A' as follows:

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

Simple, but effective.

26.3 Conditional

Consider the following argument:

> If Zora is talented then she will be successful. Zora is talented. So Zora will be successful.

This argument is certainly valid. And it suggests a straightforward conditional elimination rule ($\rightarrow$E):

$$
\begin{array}{r|ll}
m & \mathcal{A} \to \mathcal{B} & \\
n & \mathcal{A} & \\
& \mathcal{B} & \to\text{E } m,\, n
\end{array}
$$

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 '→', having started with a sentence that did contain '→'. 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 →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:

> Willie is a stoner. Therefore if Willie is a Texan, then Willie is both a stoner *and* a Texan.

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

> Assume that Willie is a stoner. Now, *additionally* assume that Willie is a Texan. Then by conjunction introduction—which we just discussed—Willie is both a stoner and a Texan. Of course, that's conditional on the assumption that Willie is a Texan. But this just means that, if Willie is a Texan, then he is both a stoner and a Texan.

Transferred into natural deduction format, here is the pattern of reasoning that we just used. We started with one premise, 'Willie is a stoner', thus:

$$
1 \quad \big|\ R
$$

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

$$
\begin{array}{r|l}
1 & R \\
2 & \quad \big|\ L
\end{array}
$$

Note that we are *not* claiming, on line 2, to have proved 'L' from line 1. So we do not need to write in any justification for the additional assumption on line 2. We do, however, need to mark that it is an additional assumption. We do this by drawing a line under it (to indicate that it is an assumption) and by indenting it with a further vertical line (to indicate that it is additional).

With this extra assumption in place, we are in a position to use ∧I. So we could continue our proof:

$$\begin{array}{ll} 1 & R \\ 2 & \quad L \\ 3 & \quad R \land L \qquad \land\text{I } 1, 2 \end{array}$$

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

$$\begin{array}{ll} 1 & R \\ 2 & \quad L \\ 3 & \quad R \land L \qquad \land\text{I } 1, 2 \\ 4 & L \to (R \land L) \qquad \to\text{I } 2\text{--}3 \end{array}$$

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

The 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:

$$\begin{array}{ll} i & \quad \mathcal{A} \\ j & \quad \mathcal{B} \\ & \mathcal{A} \to \mathcal{B} \qquad \to\text{I } i\text{--}j \end{array}$$

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 →I in action. Suppose we want to consider the following:

$$P \to Q, Q \to R \therefore P \to R$$

We start by listing *both* of our premises. Then, since we want to arrive at a conditional (namely, '$P \to R$'), we additionally assume the antecedent to that conditional. Thus our main proof starts:

$$\begin{array}{l|l}
1 & P \to Q \\
2 & Q \to R \\
3 & \quad\boxed{P} \\
\end{array}$$

Note that we have made 'P' available, by treating it as an additional assumption. But now, we can use $\to$E on the first premise. This will yield 'Q'. And we can then use $\to$E on the second premise. So, by assuming 'P' we were able to prove 'R', so we apply the $\to$I rule—discharging 'P'—and finish the proof. Putting all this together, we have:

$$\begin{array}{l|l l}
1 & P \to Q & \\
2 & Q \to R & \\
3 & \quad P & \\
4 & \quad Q & \to\text{E } 1, 3 \\
5 & \quad R & \to\text{E } 2, 4 \\
6 & P \to R & \to\text{I } 3\text{–}5 \\
\end{array}$$

26.4   Additional assumptions and subproofs

The rule $\to$I invoked the idea of making additional assumptions. These need to be handled with some care.

Consider this proof:

$$\begin{array}{l|l l}
1 & A & \\
2 & \quad B & \\
3 & \quad B \land B & \land\text{I } 2, 2 \\
4 & \quad B & \land\text{E } 3 \\
5 & B \to B & \to\text{I } 2\text{–}4 \\
\end{array}$$

This is perfectly in keeping with the rules we have laid down already. And it should not seem particularly strange. Since '$B \to B$' is a tautology, no particular premises should be required to prove it.

But suppose we now tried to continue the proof as follows:

$$
\begin{array}{ll}
1 & A \\
2 & \quad B \\
3 & \quad B \land B \qquad \land\text{I } 2,\,2 \\
4 & \quad B \qquad\qquad \land\text{E } 3 \\
5 & B \to B \qquad\qquad \to\text{I } 2\text{–}4 \\
6 & B \qquad\qquad\qquad\quad \text{naughty attempt to invoke } \to\text{E } 5,\,4
\end{array}
$$

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 Zora is talented (symbolized by 'A'), I shouldn't be able to conclude that $2 + 2 = 5$ (symbolized by 'B')! So we must be prohibited from doing this. But how are we to implement the prohibition?

We can describe the process of making an additional assumption as one of performing a *subproof*: a subsidiary proof within the main proof. When we start a subproof, we draw another vertical line to indicate that we are no longer in the main proof. Then we write in the assumption upon which the subproof will be based. A subproof can be thought of as essentially posing this question: *what could we show, if we also make this additional assumption?*

When we are working within the subproof, we can refer to the additional assumption that we made in introducing the subproof, and to anything that we obtained from our original assumptions. (After all, those original assumptions are still in effect.) But at some point, we will want to stop working with the additional assumption: we will want to return from the subproof to the main proof. To indicate that we have returned to the main proof, the vertical line for the subproof comes to an end. At this point, we say that the subproof is CLOSED. Having closed a subproof, we have set aside the additional assumption, so it will be illegitimate to draw upon anything that depends upon that additional assumption. Thus we stipulate:

> Any rule whose citation requires mentioning individual lines can mention any earlier lines, *except* for those lines which occur within a closed subproof.

This stipulation rules out the disastrous attempted proof above. The rule of $\to$E requires that we cite two individual lines from earlier in the proof. In the purported proof, above, one of these lines (namely, line 4) occurs within a subproof that has (by line 6) been closed. This is illegitimate.

Closing a subproof is called DISCHARGING the assumptions of that subproof. So we can put the point this way: *you cannot refer back to anything that was obtained using discharged assumptions.*

Subproofs, then, allow us to think about what we could show, if we made additional assumptions. The point to take away from this is not surprising—in the course of a proof, we have to keep very careful track of what assumptions we

are making, at any given moment. Our proof system does this very graphically. (Indeed, that's precisely why we have chosen to use *this* proof system.)

Once we have started thinking about what we can show by making additional assumptions, nothing stops us from posing the question of what we could show if we were to make *even more* assumptions? This might motivate us to introduce a subproof within a subproof. Here is an example which only uses the rules of proof that we have considered so far:

$$
\begin{array}{ll}
1 \quad & A \\
2 \quad & \quad B \\
3 \quad & \qquad C \\
4 \quad & \qquad A \land B \qquad \land\text{I } 1,\, 2 \\
5 \quad & \quad C \to (A \land B) \qquad \to\text{I } 3\text{--}4 \\
6 \quad & B \to (C \to (A \land B)) \qquad \to\text{I } 2\text{--}5
\end{array}
$$

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:

$$
\begin{array}{ll}
1 \quad & A \\
2 \quad & \quad B \\
3 \quad & \qquad C \\
4 \quad & \qquad A \land B \qquad \land\text{I } 1,\, 2 \\
5 \quad & \quad C \to (A \land B) \qquad \to\text{I } 3\text{--}4 \\
6 \quad & B \to (C \to (A \land B)) \qquad \to\text{I } 2\text{--}5 \\
7 \quad & C \to (A \land B) \qquad\quad \text{naughty attempt to invoke } \to\text{I } 3\text{--}4
\end{array}
$$

This would be awful. If I tell you that Zora is talented, you should not be able to infer that, if Willie is from Texas (symbolized by 'C') then *both* Zora is talented and $2 + 2 = 5$! But this is just what such a proof would suggest, if it were permissible.

The essential problem is that the subproof that began with the assumption 'C' depended crucially on the fact that we had assumed 'B' on line 2. By line 6, we have *discharged* the assumption 'B': we have stopped asking ourselves what we could show, if we also assumed 'B'. So it is simply cheating, to try to help ourselves (on line 7) to the subproof that began with the assumption 'C'. Thus we stipulate, much as before:

> Any rule whose citation requires mentioning an entire subproof can mention any earlier subproof, *except* for those subproofs which occur within some *other* closed subproof.

The attempted disastrous proof violates this stipulation. The subproof of lines 3–4 occurs within a subproof that ends on line 5. So it cannot be invoked in line 7.

It is always permissible to open a subproof with any assumption. However, there is some strategy involved in picking a useful assumption. Starting a subproof with an arbitrary, wacky assumption would just waste lines of the proof. In order to obtain a conditional by →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 '$W \leftrightarrow X$', for instance, you must be able to prove 'X' on the assumption 'W' *and* prove 'W' on the assumption 'X'. The biconditional introduction rule ($\leftrightarrow$I) therefore requires two subproofs. Schematically, the rule works like this:

$$
\begin{array}{ll}
i & \quad \mathcal{A} \\
  & \quad \overline{} \\
j & \quad \mathcal{B} \\
k & \quad \mathcal{B} \\
  & \quad \overline{} \\
l & \quad \mathcal{A} \\[4pt]
  & \mathcal{A} \leftrightarrow \mathcal{B} \qquad \leftrightarrow\text{I } i\text{–}j,\ k\text{–}l
\end{array}
$$

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:

$$\begin{array}{ll} m & \mathcal{A} \leftrightarrow \mathcal{B} \\[1ex] n & \mathcal{A} \\[1ex] & \mathcal{B} \qquad\qquad \leftrightarrow\text{E } m,\, n \end{array}$$

and equally:

$$\begin{array}{ll} m & \mathcal{A} \leftrightarrow \mathcal{B} \\[1ex] n & \mathcal{B} \\[1ex] & \mathcal{A} \qquad\qquad \leftrightarrow\text{E } m,\, n \end{array}$$

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 ↔E, we always cite the biconditional first.

26.6 Disjunction

Suppose Willie is a stoner. Then Willie is either a stoner or a Texan. After all, to say that Willie is either a stoner or a Texan is to say something weaker than to say that Willie is a stoner.

Let me emphasize this point. Suppose Willie is a stoner. It follows that Willie is *either* a stoner *or* a kumquat. Equally, it follows that *either* Willie is a stoner *or* that kumquats are the only fruit. Equally, it follows that *either* Willie is a stoner *or* that God is dead. Many of these things are strange inferences to draw. But there is nothing *logically* wrong with them (even if they maybe violate all sorts of implicit conversational norms).

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

$$\begin{array}{ll} m & \mathcal{A} \\[1ex] & \mathcal{A} \vee \mathcal{B} \qquad \vee\text{I } m \end{array}$$

and

$$\begin{array}{ll} m & \mathcal{A} \\[1ex] & \mathcal{B} \vee \mathcal{A} \qquad \vee\text{I } m \end{array}$$

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

$$
\begin{array}{ll}
1 & M \\
\hline
2 & M \lor ([(A \leftrightarrow B) \to (C \land D)] \leftrightarrow [E \land F]) \qquad \lor\text{I } 1
\end{array}
$$

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

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

But suppose that we could somehow show both of the following: first, that Willie's being a stoner entails that he is an Austrian economist: second, that Willie's being a Texan entails that he is an Austrian economist. Then if we know that Willie is either a stoner or a Texan, then we know that, whichever he is, Willie is an Austrian economist. This insight can be expressed in the following rule, which is our disjunction elimination ($\lor$E) rule:

$$
\begin{array}{ll}
m & \mathcal{A} \lor \mathcal{B} \\
i & \quad\quad \mathcal{A} \\
j & \quad\quad \mathcal{C} \\
k & \quad\quad \mathcal{B} \\
l & \quad\quad \mathcal{C} \\
& \mathcal{C} \qquad\qquad \lor\text{E } m,\ i\text{--}j,\ k\text{--}l
\end{array}
$$

This is obviously a bit clunkier to write down than our previous rules, but the point is fairly simple. Suppose we have some disjunction, $\mathcal{A} \lor \mathcal{B}$. Suppose we have two subproofs, showing us that $\mathcal{C}$ follows from the assumption that $\mathcal{A}$, and that $\mathcal{C}$ follows from the assumption that $\mathcal{B}$. Then we can infer $\mathcal{C}$ itself. As usual, there can be as many lines as you like between i and j, and as many lines as you like between k and l. Moreover, the subproofs and the disjunction can come in any order, and do not have to be adjacent.

Some examples might help illustrate this. Consider this argument:

$$(P \land Q) \lor (P \land R) \therefore P$$

An example proof might run thus:

$$
\begin{array}{l l l}
1 & (P \wedge Q) \vee (P \wedge R) & \\
2 & \quad P \wedge Q & \\
3 & \quad P & \wedge\text{E } 2 \\
4 & \quad P \wedge R & \\
5 & \quad P & \wedge\text{E } 4 \\
6 & P & \vee\text{E } 1,\ 2\text{--}3,\ 4\text{--}5
\end{array}
$$

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:

$$
\begin{array}{l l l}
1 & A \wedge (B \vee C) & \\
2 & A & \wedge\text{E } 1 \\
3 & B \vee C & \wedge\text{E } 1 \\
4 & \quad B & \\
5 & \quad A \wedge B & \wedge\text{I } 2,\ 4 \\
6 & \quad (A \wedge B) \vee (A \wedge C) & \vee\text{I } 5 \\
7 & \quad C & \\
8 & \quad A \wedge C & \wedge\text{I } 2,\ 7 \\
9 & \quad (A \wedge B) \vee (A \wedge C) & \vee\text{I } 8 \\
10 & (A \wedge B) \vee (A \wedge C) & \vee\text{E } 3,\ 4\text{--}6,\ 7\text{--}9
\end{array}
$$

Don't be alarmed if you think that you wouldn't have been able to come up with this proof yourself. The ability to come up with novel proofs will come with practice. The key question at this stage is whether, looking at the proof, you can see that it conforms with the rules that we have laid down. And that just involves checking every line, and making sure that it is justified in accordance with the rules we have laid down.

26.7   Contradiction

We have only one connective left to deal with: negation. But we will not tackle negation directly. Instead, we will first think about *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, '⊥', 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:

$$
\begin{array}{l|l}
m & \mathcal{A} \\[2ex]
n & \neg\mathcal{A} \\[2ex]
  & \bot \qquad \bot\text{I } m,\, n
\end{array}
$$

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.

Our elimination rule for '⊥' is known as *ex falso quod libet*. This means 'anything follows from a contradiction'. And the idea is precisely that: if we obtained a contradiction, symbolized by '⊥', then we can infer whatever we like. How can this be motivated, as a rule of argumentation? Well, consider the English rhetorical device '...and if *that's* true, I'll eat my hat'. Since contradictions simply cannot be true, if one *is* true then not only will I eat my hat, I'll have it too. Here is the formal rule:

$$
\begin{array}{l|l}
m & \bot \\[2ex]
  & \mathcal{A} \qquad \bot\text{E } m
\end{array}
$$

Note that $\mathcal{A}$ can be *any* sentence whatsoever.

A final remark. I have said that '⊥' 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 '⊥' as a new atomic sentence of TFL, but one which can only ever have the truth value False.
- We might regard '⊥' 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 '⊥', 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 will *officially* plump for the second. '⊥' 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   Negation

There is obviously a tight link between contradiction and negation. Indeed, the ⊥I rule essentially behaves as a rule for negation elimination: we introduce '⊥' when a sentence and its negation both appear in our proof. So there is no need for us to add a further rule for negation elimination.

However, we do need to state 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:

$$
\begin{array}{lll}
i & \quad \begin{array}{|l} \mathcal{A} \\ \hline \end{array} & \\
j & \quad \begin{array}{|l} \bot \end{array} & \\
& \neg\mathcal{A} & \quad \neg\text{I } i\text{--}j
\end{array}
$$

There can be as many lines between i and j as you like. To see this in practice, and interacting with negation, consider this proof:

$$
\begin{array}{lll}
1 & \begin{array}{|l} D \end{array} & \\
2 & \quad \begin{array}{|l} \neg D \end{array} & \\
3 & \quad \begin{array}{|l} \bot \end{array} & \quad \bot\text{I } 1,\ 2 \\
4 & \neg\neg D & \quad \neg\text{I } 2\text{--}3
\end{array}
$$

We will also add another rule for negation. It is 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 Maya will have brought an umbrella (for fear of burning). Suppose we can also show that, if it's not sunny outside, then Maya will have brought an umbrella (for fear of rain). Well, there is no third way for the weather to be. So, *whatever the weather*, Maya will have brought an umbrella.

This line of thinking motivates the following rule:

$$
\begin{array}{ll}
i & \quad\quad \mathcal{A} \\
j & \quad\quad \mathcal{B} \\
k & \quad\quad \neg\mathcal{A} \\
l & \quad\quad \mathcal{B} \\
\\
& \mathcal{B} \quad\quad\quad\quad \text{TND } i\text{--}j,\ k\text{--}l
\end{array}
$$

The rule is sometimes called *tertium non datur*, which means 'no third way'. There can be as many lines as you like between i and j, and as many lines as you like between k and l. Moreover, the subproofs can come in any order, and the second subproof does not need to come immediately after the first.

To see the rule in action, consider:

$$P \therefore (P \wedge D) \vee (P \wedge \neg D)$$

Here is a proof corresponding with the argument:

$$
\begin{array}{lll}
1 & P & \\
2 & \quad D & \\
3 & \quad P \wedge D & \wedge\text{I } 1,\ 2 \\
4 & \quad (P \wedge D) \vee (P \wedge \neg D) & \vee\text{I } 3 \\
5 & \quad \neg D & \\
6 & \quad P \wedge \neg D & \wedge\text{I } 1,\ 5 \\
7 & \quad (P \wedge D) \vee (P \wedge \neg D) & \vee\text{I } 6 \\
8 & (P \wedge D) \vee (P \wedge \neg D) & \text{TND } 2\text{--}4,\ 5\text{--}7
\end{array}
$$

These are all of the basic rules for the proof system for TFL.

Practice exercises

A. The following two 'proofs' are *incorrect*. Explain the mistakes they make.

$$
\begin{array}{lll}
1 & \neg L \rightarrow (A \wedge L) & \\
2 & \quad \neg L & \\
3 & \quad A & \rightarrow\text{E } 1,\ 2 \\
4 & \quad L & \\
5 & \quad \bot & \bot\text{I } 4,\ 2 \\
6 & \quad A & \bot\text{E } 5 \\
7 & A & \text{TND } 2\text{--}3,\ 4\text{--}6
\end{array}
$$

$$
\begin{array}{l|ll}
1 & A \wedge (B \wedge C) & \\
2 & (B \vee C) \to D & \\
\hline
3 & B & \wedge\text{E } 1 \\
4 & B \vee C & \vee\text{I } 3 \\
5 & D & \to\text{E } 4,\ 2 \\
\end{array}
$$

B. The following three proofs are missing their citations (rule and line numbers). Add them, to turn them into bona fide proofs. Additionally, write down the argument that corresponds to each proof.

$$
\begin{array}{l|l}
1 & P \wedge S \\
2 & S \to R \\
\hline
3 & P \\
4 & S \\
5 & R \\
6 & R \vee E \\
\end{array}
\qquad
\begin{array}{l|l}
1 & \neg L \to (J \vee L) \\
2 & \neg L \\
\hline
3 & J \vee L \\
4 & \quad J \\
5 & \quad J \wedge J \\
6 & \quad J \\
7 & \quad L \\
8 & \quad \bot \\
9 & \quad J \\
10 & J \\
\end{array}
$$

$$
\begin{array}{l|l}
1 & A \to D \\
\hline
2 & \quad A \wedge B \\
3 & \quad A \\
4 & \quad D \\
5 & \quad D \vee E \\
6 & (A \wedge B) \to (D \vee E) \\
\end{array}
$$

C. Give a proof for each of the following arguments:

1. $J \to \neg J \therefore \neg J$
2. $Q \to (Q \wedge \neg Q) \therefore \neg Q$
3. $A \to (B \to C) \therefore (A \wedge B) \to C$
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. $\neg F \to G, F \to H \therefore G \vee H$
8. $(Z \wedge K) \vee (K \wedge M), K \to D \therefore D$
9. $P \wedge (Q \vee R), P \to \neg R \therefore Q \vee E$
10. $S \leftrightarrow T \therefore S \leftrightarrow (T \vee S)$
11. $\neg(P \to Q) \therefore \neg Q$

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

Additional rules for TFL 27

In §26, we introduced the basic rules of our proof system for TFL. In this section, we will 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:

$$
\begin{array}{r|ll}
m & \mathcal{A} & \\
 & \mathcal{A} & \text{R } m
\end{array}
$$

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

$$
\begin{array}{r|l}
1 & A \to \neg A \\
\hline
2 & \quad A \\
3 & \quad \neg A \qquad \to\text{E } 1, 2 \\
4 & \quad \neg A \\
5 & \quad \neg A \qquad \text{R } 4 \\
6 & \neg A \qquad\quad \text{TND } 2\text{–}3, 4\text{–}5
\end{array}
$$

This is a fairly typical use of the R rule.

27.2 Disjunctive syllogism

Here is a very natural argument form.

> Miles is either at rehearsal or at the party. He is not at the party. So, he is at rehearsal.

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

$$
\begin{array}{ll}
m & \mathcal{A} \lor \mathcal{B} \\[1em]
n & \neg\mathcal{A} \\[1em]
  & \mathcal{B} \qquad\qquad \text{DS } m, n
\end{array}
$$

and

$$
\begin{array}{ll}
m & \mathcal{A} \lor \mathcal{B} \\[1em]
n & \neg\mathcal{B} \\[1em]
  & \mathcal{A} \qquad\qquad \text{DS } m, n
\end{array}
$$

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 Miles was able to reschedule rehearsal, he'd be at the party. Miles isn't at the party. So he wasn't able to reschedule rehearsal.

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

$$
\begin{array}{ll}
m & \mathcal{A} \to \mathcal{B} \\[1em]
n & \neg\mathcal{B} \\[1em]
  & \neg\mathcal{A} \qquad\qquad \text{MT } m, n
\end{array}
$$

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*. This rule does exactly what it says on the tin:

$$
\begin{array}{r|l}
m & \neg\neg\mathcal{A} \\[2ex]
 & \mathcal{A} \qquad\qquad \text{DNE } m
\end{array}
$$

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: 'Maya is not *not* happy'. Arguably, one cannot infer 'Maya is happy', since the first sentence should be understood as meaning the same as 'Maya is not *un*happy'. This is compatible with 'Maya 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. (These are named after August De Morgan.) The shape of the rules should be familiar from truth tables.

The first De Morgan rule is:

$$
\begin{array}{r|l}
m & \neg(\mathcal{A} \land \mathcal{B}) \\[2ex]
 & \neg\mathcal{A} \lor \neg\mathcal{B} \qquad \text{DeM } m
\end{array}
$$

The second De Morgan is the reverse of the first:

$$
\begin{array}{r|l}
m & \neg\mathcal{A} \lor \neg\mathcal{B} \\[2ex]
 & \neg(\mathcal{A} \land \mathcal{B}) \qquad \text{DeM } m
\end{array}
$$

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

$$
\begin{array}{r|l}
m & \neg(\mathcal{A} \lor \mathcal{B}) \\[2ex]
 & \neg\mathcal{A} \land \neg\mathcal{B} \qquad \text{DeM } m
\end{array}
$$

And the fourth is the reverse of the third:

$$
\begin{array}{r|l}
m & \neg\mathcal{A} \land \neg\mathcal{B} \\[2ex]
 & \neg(\mathcal{A} \lor \mathcal{B}) \qquad \text{DeM } m
\end{array}
$$

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	$\bot$
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	$\bot$
11	$\neg Z$
12	$N \wedge \neg C$
13	N
14	$\bot$
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. $(X \wedge Y) \vee (X \wedge Z)$, $\neg(X \wedge D)$, $D \vee M$ $\therefore M$

We will introduce some new vocabulary. The following expression:

$$\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \vdash \mathcal{C}$$

means that there is some proof which starts with assumptions among $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ and ends with $\mathcal{C}$ (and no undischarged assumptions other than those we started with). Derivatively, we will write:

$$\vdash \mathcal{A}$$

to mean that there is a proof of $\mathcal{A}$ with no assumptions.

The symbol '$\vdash$' is called the *single turnstile*. I want to emphasize that this is not the double turnstile symbol ('$\vDash$') that we used to symbolize 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 very different notions.*

Armed with our '$\vdash$' symbol, we can introduce a new terminology.

> $\mathcal{A}$ is a THEOREM iff $\vdash \mathcal{A}$

To illustrate this, suppose I want to prove that '$\neg(A \wedge \neg A)$' is a theorem. So I must start my proof without *any* assumptions. However, since I want to prove a sentence whose main logical operator is a negation, I will want to immediately begin a subproof, with the additional assumption '$A \wedge \neg A$', and show that this 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	$\bot$	$\bot$I 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 $\mathcal{A}$ and $\mathcal{B}$ are PROVABLY EQUIVALENT iff each can be proved from the other; i.e., both $\mathcal{A} \vdash \mathcal{B}$ and $\mathcal{B} \vdash \mathcal{A}$.

As in the case of showing that a sentence is a theorem, it is relatively easy to show that two sentences are provably equivalent: it just requires a pair of proofs. Showing that sentences are *not* provably equivalent would be much harder: it is just as hard as showing that a sentence is not a theorem.

Here is a third, related, bit of terminology:

> The sentences $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ are JOINTLY CONTRARY iff a contradiction can be proved from them, i.e. $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \vdash \bot$.

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 summarizes whether one or two proofs suffice, or whether we must reason about all possible proofs.

	Yes	**No**
theorem?	one proof	all possible proofs
equivalent?	two proofs	all possible proofs
not contrary?	all possible proofs	one proof

Practice exercises

A. Show that each of the following sentences is a theorem:

1. $O \to O$
2. $N \vee \neg N$
3. $J \leftrightarrow [J \vee (L \wedge \neg L)]$
4. $((A \to B) \to A) \to A$

B. Provide proofs to show each of the following:

1. $C \to (E \wedge G), \neg C \to G \vdash G$
2. $M \wedge (\neg N \to \neg M) \vdash (N \wedge M) \vee \neg M$
3. $(Z \wedge K) \leftrightarrow (Y \wedge M), D \wedge (D \to M) \vdash Y \to Z$

    4. $(W \lor X) \lor (Y \lor Z), X \to Y, \neg Z \vdash W \lor Y$

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. $T \to S$, $\neg S \to \neg T$
    4. $U \to I$, $\neg(U \land \neg I)$
    5. $\neg(C \to D)$, $C \land \neg D$
    6. $\neg G \leftrightarrow H$, $\neg(G \leftrightarrow H)$

D. If you know that $\mathcal{A} \vdash \mathcal{B}$, what can you say about $(\mathcal{A} \land \mathcal{C}) \vdash \mathcal{B}$? What about $(\mathcal{A} \lor \mathcal{C}) \vdash \mathcal{B}$? Explain your answers.

E. In this section, I claimed that it is just as hard to show that two sentences are not provably equivalent, as it is to show that a sentence is not a theorem. Why did I claim this? (*Hint*: think of a sentence that would be a theorem iff $\mathcal{A}$ and $\mathcal{B}$ were provably equivalent.)

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. Ask what you could do to get to this new goal.

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

Work forwards from what you have. When you are starting a proof, look at the premises; later, look at the sentences that you have obtained so far. Think about the elimination rules for the main operators of these sentences. These will tell you what your options are.

    For a short proof, you might be able to eliminate the premises and introduce the conclusion. A long proof is formally just a number of short proofs linked together, so you can fill the gap by alternately working back from the conclusion and forward from the premises.

Try proceeding indirectly. If you cannot find a way to show $\mathcal{A}$ directly, try starting by assuming $\neg\mathcal{A}$. If a contradiction follows, then you will be able to obtain $\neg\neg\mathcal{A}$ by $\neg$I, and then $\mathcal{A}$ by DNE.

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

Derived rules 30

In this section, we will see 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:

$$m \quad \Big| \quad \mathcal{A}$$

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:

$$
\begin{array}{ll}
m & \mathcal{A} \\
k & \mathcal{A} \wedge \mathcal{A} \quad \wedge\text{I } m \\
k+1 & \mathcal{A} \quad\quad\quad \wedge\text{E } k
\end{array}
$$

To be clear: this is not a proof. Rather, it is a proof *scheme*. After all, it uses a variable, '$\mathcal{A}$', rather than a sentence of TFL. But the point is simple. Whatever sentences of TFL we plugged in for '$\mathcal{A}$', 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 that, 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 the rule R as a *derived* rule, since its justification is derived from our basic rules.

30.2 Derivation of Disjunctive syllogism

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

$$
\begin{array}{c|l}
m & \mathcal{A} \vee \mathcal{B} \\[4pt]
n & \neg\mathcal{A}
\end{array}
$$

You now want, on line k, to prove $\mathcal{B}$. 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:

$$
\begin{array}{c|ll}
m & \mathcal{A} \vee \mathcal{B} & \\[4pt]
n & \neg\mathcal{A} & \\[4pt]
k & \quad \mid \mathcal{A} & \\[4pt]
k+1 & \quad \mid \bot & \bot\text{I } k,\, n \\[4pt]
k+2 & \quad \mid \mathcal{B} & \bot\text{E } k+1 \\[4pt]
k+3 & \quad \mid \mathcal{B} & \\[4pt]
k+4 & \quad \mid \mathcal{B} \wedge \mathcal{B} & \wedge\text{I } k+3,\, k+3 \\[4pt]
k+5 & \quad \mid \mathcal{B} & \wedge\text{E } k+4 \\[4pt]
k+6 & \mathcal{B} & \vee\text{E } m,\, k\text{--}k+2,\, k+3\text{--}k+5
\end{array}
$$

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 take a few extra lines and prove the same thing using only our basic rules. It is a *derived* rule.

30.3   Derivation of Modus Tollens

Suppose you have the following in your proof:

$$
\begin{array}{c|l}
m & \mathcal{A} \rightarrow \mathcal{B} \\[4pt]
n & \neg\mathcal{B}
\end{array}
$$

You now want, on line k, to prove $\neg\mathcal{A}$. 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:

$$
\begin{array}{c|ll}
m & \mathcal{A} \rightarrow \mathcal{B} & \\[4pt]
n & \neg\mathcal{B} & \\[4pt]
k & \quad \mid \mathcal{A} & \\[4pt]
k+1 & \quad \mid \mathcal{B} & \rightarrow\text{E } m,\, k \\[4pt]
k+2 & \quad \mid \bot & \bot\text{I } k+1,\, n \\[4pt]
k+3 & \neg\mathcal{A} & \neg\text{I } k\text{--}k+2
\end{array}
$$

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

30.4 Derivation of Double-negation elimination

Consider the following deduction scheme:

$$
\begin{array}{lll}
m & \neg\neg\mathcal{A} & \\
k & \quad \mathcal{A} & \\
k+1 & \quad \mathcal{A} & \text{R } k \\
k+2 & \quad \neg\mathcal{A} & \\
k+3 & \quad \bot & \bot\text{I } k+2,\, m \\
k+4 & \quad \mathcal{A} & \bot\text{E } k+3 \\
k+5 & \mathcal{A} & \text{TND } k\text{–}k+1,\, k+2\text{–}k+4
\end{array}
$$

Again, then, we can derive the DNE rule from the *basic* rules of §26.

30.5 Derivation of De Morgan rules

Here is a demonstration of how we could derive the first De Morgan rule:

$$
\begin{array}{lll}
m & \neg(\mathcal{A} \wedge \mathcal{B}) & \\
k & \quad \mathcal{A} & \\
k+1 & \qquad \mathcal{B} & \\
k+2 & \qquad \mathcal{A} \wedge \mathcal{B} & \wedge\text{I } k,\, k+1 \\
k+3 & \qquad \bot & \bot\text{I } k+2,\, m \\
k+4 & \quad \neg\mathcal{B} & \neg\text{I } k+1\text{–}k+3 \\
k+5 & \quad \neg\mathcal{A} \vee \neg\mathcal{B} & \vee\text{I } k+4 \\
k+6 & \quad \neg\mathcal{A} & \\
k+7 & \quad \neg\mathcal{A} \vee \neg\mathcal{B} & \vee\text{I } k+6 \\
k+8 & \neg\mathcal{A} \vee \neg\mathcal{B} & \text{TND } k\text{–}k+5,\, k+6\text{–}k+7
\end{array}
$$

Here is a demonstration of how we could derive the second De Morgan rule:

m	$\neg\mathcal{A} \vee \neg\mathcal{B}$	
k	$\mathcal{A} \wedge \mathcal{B}$	
$k+1$	$\mathcal{A}$	$\wedge$E k
$k+2$	$\mathcal{B}$	$\wedge$E k
$k+3$	$\neg\mathcal{A}$	
$k+4$	$\bot$	$\bot$I $k+1$, $k+3$
$k+5$	$\neg\mathcal{B}$	
$k+6$	$\bot$	$\bot$I $k+2$, $k+5$
$k+7$	$\bot$	$\vee$E m, $k+3$–$k+4$, $k+5$–$k+6$
$k+8$	$\neg(\mathcal{A} \wedge \mathcal{B})$	$\neg$I k–$k+7$

Similar demonstrations can be offered explaining how we could derive the third and fourth De Morgan rules. These are left as exercises.

Practice exercises

A. 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.)

Chapter 7

Natural deduction for FOL

Basic rules for FOL   31

FOL makes use of all of the connectives of TFL. So proofs in FOL will use all of the basic and derived rules from chapter 6. We will also use the proof-theoretic notions (particularly, the symbol '⊢') 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:

$$
\begin{array}{l|l}
1 & \forall x Rxxd \\
\hline
2 & Raad \qquad \forall \text{E } 1
\end{array}
$$

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

$$
\begin{array}{l|l}
1 & \forall x Rxxd \\
\hline
2 & Rddd \qquad \forall \text{E } 1
\end{array}
$$

We obtained line 2 here by dropping the universal quantifier and replacing every instance of 'x' with 'd'. We could have done the same with any other name we wanted.

This motivates the universal elimination rule (∀E):

$$
\begin{array}{l|l}
m & \forall x \mathcal{A}(\ldots x \ldots x \ldots) \\
 & \mathcal{A}(\ldots c \ldots c \ldots) \qquad \forall \text{E } m
\end{array}
$$

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 emphasize that (as with every elimination rule) you can only apply the ∀E rule when the universal quantifier is the main logical operator. Thus the following is outright banned:

$$1 \quad \mid \quad \forall x Bx \to Bk$$

$$2 \quad \mid \quad Bb \to Bk \qquad \text{naughtily attempting to invoke } \forall E\ 1$$

This is illegitimate, since '∀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:

$$1 \quad \mid \quad Raad$$

$$2 \quad \mid \quad \exists x Raax \qquad \exists I\ 1$$

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

$$1 \quad \mid \quad Raad$$

$$2 \quad \mid \quad \exists x Rxxd \qquad \exists I\ 1$$

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. (After all, if Narcissus loves himself, then there is someone who loves Narcissus.) So we would also allow:

$$1 \quad \mid \quad Raad$$

$$2 \quad \mid \quad \exists x Rxad \qquad \exists I\ 1$$

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 will need to introduce some new notation.

Where $\mathcal{A}$ is a sentence containing the name c, we can emphasize this by writing '$\mathcal{A}(\ldots c \ldots c \ldots)$'. We will write '$\mathcal{A}(\ldots \chi \ldots c \ldots)$' to indicate any formula obtained by replacing *some or all* of the instances of the name c with the variable χ. Armed with this, our introduction rule is:

$$m \quad \mid \quad \mathcal{A}(\ldots c \ldots c \ldots)$$

$$ \quad \mid \quad \exists \chi \mathcal{A}(\ldots \chi \ldots c \ldots) \qquad \exists I\ m$$

χ must not occur in $\mathcal{A}(\ldots c \ldots c \ldots)$

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

1	$Raad$	
2	$\exists x Rxad$	$\exists$I 1
3	$\exists y \exists x Rxyd$	$\exists$I 2

But this is banned:

1	$Raad$	
2	$\exists x Rxad$	$\exists$I 1
3	$\exists x \exists x Rxxd$	naughtily attempting to invoke $\exists$I 2

since the expression on line 3 contains clashing variables, and so fails to count as a sentence of FOL.

31.3   Empty domains

The following proof combines our two new rules for quantifiers:

1	$\forall x Fx$	
2	Fa	$\forall$E 1
3	$\exists x Fx$	$\exists$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 following that something is F, for there is nothing to *be* F. So if we claim that, as a matter of logic alone, '$\exists x Fx$' follows from '$\forall x Fx$', 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:

- $\forall x Fx \vdash Fa$: 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 will 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:

$$
\begin{array}{c|ll}
1 & \forall xFx & \\
\hline
2 & Fa & \forall \text{E } 1
\end{array}
$$

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 '$\forall y Fy$', thus:

$$
\begin{array}{ll}
1 & \forall x Fx \\
2 & Fa \qquad \forall \text{E } 1 \\
3 & \forall y Fy \qquad \forall \text{I } 2
\end{array}
$$

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):

$$
\begin{array}{ll}
m & \mathcal{A}(\ldots c \ldots c \ldots) \\
 & \forall \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots) \qquad \forall \text{I } m
\end{array}
$$

c must not occur in any undischarged assumption
χ must not occur in $\mathcal{A}(\ldots c \ldots c \ldots)$

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.[1] To see the constraint in action, consider this terrible argument:

Everyone loves Beyoncé; therefore everyone loves themselves.

We might symbolize this obviously invalid inference pattern as:

$$\forall x Lxk \therefore \forall x Lxx$$

Now, suppose we tried to offer a proof that vindicates this argument:

$$
\begin{array}{ll}
1 & \forall x Lxk \\
2 & Lkk \qquad \forall \text{E } 1 \\
3 & \forall x Lxx \qquad \text{naughtily attempting to invoke } \forall \text{I } 2
\end{array}
$$

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.

---

[1]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 will treat '$\perp$' as a canonical contradiction *that involves no particular names*.

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

$$
\begin{array}{lll}
1 & \quad\begin{array}{|l} Gd \\ \hline \end{array} & \\
2 & \quad\begin{array}{|l} Gd \end{array} & \text{R } 1 \\
3 & Gd \rightarrow Gd & \rightarrow\text{I } 1\text{--}2 \\
4 & \forall z(Gz \rightarrow Gz) & \forall\text{I } 3
\end{array}
$$

This tells us that '$\forall z(Gz \rightarrow Gz)$' is a *theorem*. And that is as it should be.

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:

$$
\begin{array}{lll}
1 & \exists xFx & \\
2 & \forall x(Fx \rightarrow Gx) & \\
3 & \quad\begin{array}{|l} Fo \\ \hline \end{array} & \\
4 & \quad\begin{array}{|l} Fo \rightarrow Go \end{array} & \forall\text{E } 2 \\
5 & \quad\begin{array}{|l} Go \end{array} & \rightarrow\text{E } 4,\, 3 \\
6 & \quad\begin{array}{|l} \exists xGx \end{array} & \exists\text{I } 5 \\
7 & \exists xGx & \exists\text{E } 1,\, 3\text{--}6
\end{array}
$$

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 '$\exists x Gx$' on its own.

Putting this together, we obtain the existential elimination rule ($\exists$E):

$$
\begin{array}{ll}
m & \quad \exists x \mathcal{A}(\ldots x \ldots x \ldots) \\[6pt]
i & \qquad\quad \mathcal{A}(\ldots c \ldots c \ldots) \\[6pt]
j & \qquad\quad \mathcal{B} \\[6pt]
& \quad \mathcal{B} \qquad\qquad\qquad \exists \text{E } m, i\text{–}j
\end{array}
$$

c must not occur in any assumption undischarged before line i
c must not occur in $\exists x \mathcal{A}(\ldots x \ldots x \ldots)$
c must not occur in $\mathcal{B}$

As with universal introduction, the constraints are extremely important. To see why, consider the following terrible argument:

> Kathryn Lindeman is a professor. There is someone who is not a professor. So Kathryn Lindeman is both a professor and not a professor.

We might symbolize 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:

$$
\begin{array}{ll}
1 & \quad Lb \\[4pt]
2 & \quad \exists x \neg Lx \\[4pt]
3 & \qquad\quad \neg Lb \\[4pt]
4 & \qquad\quad Lb \wedge \neg Lb \qquad \wedge \text{E } 1, 3 \\[4pt]
5 & \quad Lb \wedge \neg Lb \qquad\qquad \text{naughtily attempting to invoke } \exists \text{E } 2, 3\text{–}4
\end{array}
$$

The last line of the proof is not allowed. The name that we used in our substitution instance for '$\exists x \neg Lx$' on line 3, namely 'b', occurs in line 4. And the following proof would be no better:

1 | Lb
2 | $\exists x \neg Lx$
3 | $\neg Lb$
4 | $Lb \wedge \neg Lb$ $\wedge$E 1, 3
5 | $\exists x(Lx \wedge \neg Lx)$ $\exists$I 4
6 | $\exists x(Lx \wedge \neg Lx)$ naughtily attempting to invoke $\exists$E 2, 3–5

The last line of the proof would still not be allowed. For the name that we used in our substitution instance for '$\exists x \neg Lx$', namely 'b', occurs in an undischarged assumption, namely line 1.

    The moral of the story is this. *If you want to squeeze information out of an 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. The following two 'proofs' are *incorrect*. Explain why both are incorrect. Also, provide interpretations which would invalidate the fallacious argument forms the 'proofs' enshrine:

1 | $\forall x Rxx$ | | 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 | Raa |
4 | $\forall x \forall y Rxy$ | $\forall$I 3 | 4 | $\exists x Rxx$ | $\exists$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.

$$1 \quad \forall x \exists y (Rxy \lor Ryx)$$

$$2 \quad \forall x \neg Rmx$$

$$3 \quad \exists y (Rmy \lor Rym)$$

$$4 \qquad Rma \lor Ram$$

$$5 \qquad \neg Rma$$

$$6 \qquad Ram$$

$$7 \qquad \exists x Rxm$$

$$8 \quad \exists x Rxm$$

1	$\forall x (\exists y Lxy \to \forall z Lzx)$
2	Lab
3	$\exists y Lay \to \forall z Lza$
4	$\exists y Lay$
5	$\forall z Lza$
6	Lca
7	$\exists y Lcy \to \forall z Lzc$
8	$\exists y Lcy$
9	$\forall z Lzc$
10	Lcc
11	$\forall x Lxx$

1	$\forall x (Jx \to Kx)$
2	$\exists x \forall y Lxy$
3	$\forall x Jx$
4	$\forall y Lay$
5	Laa
6	Ja
7	$Ja \to Ka$
8	Ka
9	$Ka \land Laa$
10	$\exists x (Kx \land Lxx)$
11	$\exists x (Kx \land Lxx)$

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 symbolize (for example) 'No F is G' as '$\forall x (Fx \to \neg Gx)$'.

D. Aristotle and his successors identified other syllogistic forms which depended upon 'existential import'. symbolize each of the following 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. $\vdash \forall x Fx \lor \neg \forall x Fx$
2. $\vdash \forall z(Pz \lor \neg Pz)$
3. $\forall x(Ax \to Bx), \exists x Ax \vdash \exists x Bx$
4. $\forall x(Mx \leftrightarrow Nx), Ma \land \exists x Rxa \vdash \exists x Nx$
5. $\forall x \forall y Gxy \vdash \exists x Gxx$
6. $\vdash \forall x Rxx \to \exists x \exists y Rxy$
7. $\vdash \forall y \exists x(Qy \to Qx)$
8. $Na \to \forall x(Mx \leftrightarrow Ma), Ma, \neg Mb \vdash \neg Na$
9. $\forall x \forall y(Gxy \to Gyx) \vdash \forall x \forall y(Gxy \leftrightarrow Gyx)$
10. $\forall x(\neg Mx \lor Ljx), \forall x(Bx \to Ljx), \forall x(Mx \lor Bx) \vdash \forall x Ljx$

F. Write a symbolization key for the following argument, symbolize 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. $\forall x Px \to Qc, \forall x(Px \to Qc)$
2. $\forall x \forall y \forall z Bxyz, \forall x Bxxx$
3. $\forall x \forall y Dxy, \forall y \forall x Dxy$
4. $\exists x \forall y Dxy, \forall y \exists x Dxy$
5. $\forall x(Rca \leftrightarrow Rxa), Rca \leftrightarrow \forall x Rxa$

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. $\exists y \forall x Rxy \; \therefore \; \forall x \exists y Rxy$
2. $\exists x(Px \land \neg Qx) \; \therefore \; \forall x(Px \to \neg Qx)$
3. $\forall x(Sx \to Ta), Sd \; \therefore \; Ta$
4. $\forall x(Ax \to Bx), \forall x(Bx \to Cx) \; \therefore \; \forall x(Ax \to Cx)$
5. $\exists x(Dx \lor Ex), \forall x(Dx \to Fx) \; \therefore \; \exists x(Dx \land Fx)$
6. $\forall x \forall y(Rxy \lor Ryx) \; \therefore \; Rjj$
7. $\exists x \exists y(Rxy \lor Ryx) \; \therefore \; Rjj$
8. $\forall x Px \to \forall x Qx, \exists x \neg Px \; \therefore \; \exists x \neg Qx$

Conversion of quantifiers 32

In this section, we will 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 $\neg\exists x\mathcal{A}$ is logically equivalent to $\forall x\neg\mathcal{A}$. We will add some rules to our proof system that govern this. In particular, we add:

$$
\begin{array}{r|ll}
m & \forall x\neg\mathcal{A} & \\
& \neg\exists x\mathcal{A} & \text{CQ } m
\end{array}
$$

and

$$
\begin{array}{r|ll}
m & \neg\exists x\mathcal{A} & \\
& \forall x\neg\mathcal{A} & \text{CQ } m
\end{array}
$$

Equally, we add:

$$
\begin{array}{r|ll}
m & \exists x\neg\mathcal{A} & \\
& \neg\forall x\mathcal{A} & \text{CQ } m
\end{array}
$$

and

$$
\begin{array}{r|ll}
m & \neg\forall x\mathcal{A} & \\
& \exists x\neg\mathcal{A} & \text{CQ } m
\end{array}
$$

Practice exercises

A. Show that the following are jointly contrary:

1. $Sa \to Tm, Tm \to Sa, Tm \land \neg Sa$
2. $\neg \exists x Rxa, \forall x \forall y Ryx$
3. $\neg \exists x \exists y Lxy, Laa$
4. $\forall x(Px \to Qx), \forall z(Pz \to Rz), \forall y Py, \neg Qa \land \neg Rb$

B. Show that each pair of sentences is provably equivalent:

1. $\forall x(Ax \to \neg Bx), \neg \exists x(Ax \land Bx)$
2. $\forall x(\neg Ax \to Bd), \forall x Ax \lor Bd$

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. $\forall x(Fx \land Ga), \forall x Fx \land Ga$
2. $\exists x(Fx \lor Ga), \exists x Fx \lor Ga$
3. $\forall x(Ga \to Fx), Ga \to \forall x Fx$
4. $\forall x(Fx \to Ga), \exists x Fx \to Ga$
5. $\exists x(Ga \to Fx), Ga \to \exists x Fx$
6. $\exists x(Fx \to Ga), \forall x Fx \to Ga$

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 33

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 consequence of this, for our proof system, is that there are no sentences that do not already contain the identity predicate that could justify the conclusion '$a = b$'. This means that the identity introduction rule will not justify '$a = b$', or any other identity claim containing two different names.

However, every object is identical to itself. No premises, then, are required in order to conclude that something is identical to itself. So this will be the identity introduction rule:

$$c = c \qquad =\text{I}$$

Notice that this rule does not require referring to any prior lines of the proof. For any name c, you can write $c = c$ 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:

$$
\begin{array}{ll}
m & a = b \\
n & \mathcal{A}(\ldots a \ldots a \ldots) \\
  & \mathcal{A}(\ldots b \ldots a \ldots) \qquad =\text{E } m, n
\end{array}
$$

The notation here is as for $\exists$I. So $\mathcal{A}(\ldots \boldsymbol{a} \ldots \boldsymbol{a} \ldots)$ is a formula containing the name $\boldsymbol{a}$, and $\mathcal{A}(\ldots \boldsymbol{b} \ldots \boldsymbol{a} \ldots)$ is a formula obtained by replacing one or more instances of the name $\boldsymbol{a}$ with the name $\boldsymbol{b}$. 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:

$$
\begin{array}{ll}
m & \boldsymbol{a} = \boldsymbol{b} \\
n & \mathcal{A}(\ldots \boldsymbol{b} \ldots \boldsymbol{b} \ldots) \\
  & \mathcal{A}(\ldots \boldsymbol{a} \ldots \boldsymbol{b} \ldots) \qquad =\text{E } m,\, n
\end{array}
$$

This rule is sometimes called *Leibniz's Law*, after Gottfried Leibniz.

To see the rules in action, we will prove some quick results. First, we will prove that identity is *symmetric*:

$$
\begin{array}{lll}
1 & \quad a = b & \\
2 & \quad\quad a = a & =\text{I} \\
3 & \quad\quad b = a & =\text{E } 1,\, 2 \\
4 & a = b \rightarrow b = a & \rightarrow\text{I } 1\text{--}3 \\
5 & \forall y(a = y \rightarrow y = a) & \forall\text{I } 4 \\
6 & \forall x \forall y(x = y \rightarrow y = x) & \forall\text{I } 5
\end{array}
$$

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 will prove that identity is *transitive*:

$$
\begin{array}{lll}
1 & \quad a = b \wedge b = c & \\
2 & \quad\quad a = b & \wedge\text{E } 1 \\
3 & \quad\quad b = c & \wedge\text{E } 1 \\
4 & \quad\quad a = c & =\text{E } 2,\, 3 \\
5 & (a = b \wedge b = c) \rightarrow a = c & \rightarrow\text{I } 1\text{--}4 \\
6 & \forall z((a = b \wedge b = z) \rightarrow a = z) & \forall\text{I } 5 \\
7 & \forall y \forall z((a = y \wedge y = z) \rightarrow a = z) & \forall\text{I } 6 \\
8 & \forall x \forall y \forall z((x = y \wedge y = z) \rightarrow x = z) & \forall\text{I } 7
\end{array}
$$

We obtain line 4 by replacing 'b' in line 3 with 'a'; this is justified given '$a = b$'.

Derived rules 34

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:

$$
\begin{array}{lll}
1 & \forall x \neg Ax & \\
2 & \quad \exists x Ax & \\
3 & \quad\quad Ac & \\
4 & \quad\quad \neg Ac & \forall E\ 1 \\
5 & \quad\quad \bot & \bot I\ 3,\ 4 \\
6 & \quad\quad \bot & \exists E\ 2,\ 3\text{--}5 \\
7 & \neg \exists x Ax & \neg I\ 2\text{--}6
\end{array}
$$

Here is a justification of the second CQ rule:

$$
\begin{array}{lll}
1 & \exists x \neg Ax & \\
2 & \quad \forall x Ax & \\
3 & \quad\quad \neg Ac & \\
4 & \quad\quad Ac & \forall E\ 2 \\
5 & \quad\quad \bot & \bot I\ 4,\ 3 \\
6 & \quad\quad \bot & \exists E\ 1,\ 3\text{--}5 \\
7 & \neg \forall x Ax & \neg I\ 2\text{--}6
\end{array}
$$

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. Provide a proof of each claim.

1. $Pa \lor Qb, Qb \to b = c, \neg Pa \vdash Qc$
2. $m = n \lor n = o, An \vdash Am \lor Ao$
3. $\forall x\, x = m, Rma \vdash \exists x Rxx$
4. $\forall x \forall y (Rxy \to x = y) \vdash Rab \to Rba$
5. $\neg \exists x \neg x = m \vdash \forall x \forall y (Px \to Py)$
6. $\exists x Jx, \exists x \neg Jx \vdash \exists x \exists y\ \neg x = y$
7. $\forall x (x = n \leftrightarrow Mx), \forall x (Ox \lor \neg Mx) \vdash On$
8. $\exists x Dx, \forall x (x = p \leftrightarrow Dx) \vdash Dp$
9. $\exists x \big[(Kx \land \forall y (Ky \to x = y)) \land Bx \big], Kd \vdash Bd$
10. $\vdash Pa \to \forall x (Px \lor \neg x = a)$

B. Show that the following are provably equivalent:

- $\exists x \big([Fx \land \forall y (Fy \to x = y)] \land x = n \big)$
- $Fn \land \forall y (Fy \to n = y)$

And hence that both have a decent claim to symbolize the English sentence 'Nick is the F'.

C. In §17, I claimed that the following are logically equivalent symbolizations of the English sentence 'there is exactly one F':

- $\exists x Fx \land \forall x \forall y \big[(Fx \land Fy) \to x = y \big]$
- $\exists x \big[Fx \land \forall y (Fy \to x = y) \big]$
- $\exists x \forall y (Fy \leftrightarrow x = y)$

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. Symbolize 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.

E. 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:

$$\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \vdash \mathcal{C}$$

means that there is some proof which starts with assumptions $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ and ends with $\mathcal{C}$ (and no undischarged assumptions other than $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$). This is a *proof-theoretic notion*.

By contrast, this:

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

means that there is no valuation (or interpretation) which makes all of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ true and makes $\mathcal{C}$ false. This concerns assignments of truth and falsity to sentences. It is a *semantic notion*.

I cannot emphasize enough that these are different notions. But I can emphasize it a bit more: *They are different notions.*

Once you have internalized 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 will 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.

On the other hand, to show that a sentence is *not* a theorem is hard. We would need to reason about all (possible) proofs. That is very difficult. But to show that a sentence is not a logical truth, you need only construct an interpretation in which the sentence is false. Granted, it may be hard to come up with the interpretation; but once you have done so, it is relatively straightforward to check what truth value it assigns to a sentence. Given a choice between showing that a sentence is not a theorem and showing that it is not a logical truth, it would be easier to show that it is not a logical truth.

Fortunately, *a sentence is a theorem if and only if it is a logical truth.* As a result, if we provide a proof of $\mathcal{A}$ on no assumptions, and thus show that $\mathcal{A}$ is a theorem, we can legitimately infer that $\mathcal{A}$ is a logical truth; i.e., $\vDash \mathcal{A}$. Similarly, if we construct a model in which $\mathcal{A}$ is false and thus show that it is not a logical truth, it follows that $\mathcal{A}$ is not a theorem.

More generally, we have the following powerful result:

$$\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \vdash \mathcal{B} \textbf{ iff } \mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \vDash \mathcal{B}$$

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 summarizes 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 '$\vDash$' 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.[1]

---

[1] Agreement, in the case of TFL, is covered in the sequel to this book, `Metatheory`.

	Yes	No
Is $\mathcal{A}$ a **logical truth**?	give a proof which shows $\vdash \mathcal{A}$	give an interpretation in which $\mathcal{A}$ is false
Is $\mathcal{A}$ a **contradiction**?	give a proof which shows $\vdash \neg\mathcal{A}$	give an interpretation in which $\mathcal{A}$ is true
Are $\mathcal{A}$ and $\mathcal{B}$ **equivalent**?	give two proofs, one for $\mathcal{A} \vdash \mathcal{B}$ and one for $\mathcal{B} \vdash \mathcal{A}$	give an interpretation in which $\mathcal{A}$ and $\mathcal{B}$ have different truth values
Are $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ **jointly consistent**?	give an interpretation in which all of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ are true	prove a contradiction from assumptions $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$
Is $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n \therefore \mathcal{C}$ **valid**?	give a proof with assumptions $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ and concluding with $\mathcal{C}$	give an interpretation in which each of $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n$ is true and $\mathcal{C}$ is false

Appendices

Symbolic notation

<div style="text-align: right">A</div>

A.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 emphasize 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'.

A.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 recognize them if you encounter them in an article or in another book.

Negation. Two commonly used symbols are the *hoe*, '$\neg$', and the *swung dash*, '$\sim$.' In some more advanced formal systems it is necessary to distinguish between two kinds of negation; the distinction is sometimes represented by using both '$\neg$' and '$\sim$'. 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; thus an italic ampersand might appear as '*&*'.) Using this symbol is not recommended, since it is commonly used in natural English writing (e.g. 'Smith & Sons'). As a symbol in a formal system, the ampersand is not the English word '&', so it is much neater to use a completely different symbol. The most common choice now is '$\wedge$', which is a counterpart to the symbol used for disjunction. Sometimes a single dot, '$\bullet$', is used. In some older texts, there is no symbol for conjunction at all; 'A and B' is simply written 'AB'.

Material Conditional. There are two common symbols for the material conditional: the *arrow*, '$\rightarrow$', and the *hook*, '$\supset$'.

Material Biconditional. The *double-headed arrow*, '$\leftrightarrow$', is used in systems that use the arrow to represent the material conditional. Systems that use the hook for the conditional typically use the *triple bar*, '$\equiv$', for the biconditional.

Quantifiers. The universal quantifier is typically symbolized as 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 '$\forall x Px$'.

These alternative typographies are summarized below:

negation	$\neg$, $\sim$
conjunction	$\wedge$, $\&$, $\bullet$
disjunction	$\vee$
conditional	$\rightarrow$, $\supset$
biconditional	$\leftrightarrow$, $\equiv$
universal quantifier	$\forall x$, (x)

Alternative proof systems B

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.

B.1 Alternative disjunction elimination

Some systems take DS as their basic rule for disjunction elimination. Such systems can then treat the ∨E rule as a derived rule. For they might offer the following proof scheme:

m	$\mathcal{A} \lor \mathcal{B}$	
i	$\mathcal{A}$	
j	$\mathcal{C}$	
k	$\mathcal{B}$	
l	$\mathcal{C}$	
n	$\mathcal{A} \to \mathcal{C}$	$\to$I i–j
$n+1$	$\mathcal{B} \to \mathcal{C}$	$\to$I k–l
$n+2$	$\mathcal{C}$	
$n+3$	$\mathcal{C}$	R $n+2$
$n+4$	$\neg\mathcal{C}$	
$n+5$	$\mathcal{A}$	
$n+6$	$\mathcal{C}$	$\to$E n, $n+5$
$n+7$	$\bot$	$\bot$I $n+6$, $n+4$
$n+8$	$\neg\mathcal{A}$	$\neg$I $n+5$–$n+7$
$n+9$	$\mathcal{B}$	DS m, $n+8$
$n+10$	$\mathcal{C}$	$\to$E $n+1$, $n+9$
$n+11$	$\mathcal{C}$	TND $n+2$–$n+3$, $n+4$–$n+10$

So why did I choose to take $\lor$E as basic, rather than DS?[1] 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.

B.2 Alternative negation rules

Some systems take the following rule as their basic negation introduction rule:

m	$\mathcal{A}$	
$n-1$	$\mathcal{B}$	
n	$\neg\mathcal{B}$	
	$\neg\mathcal{A}$	$\neg$I* m–n

[1]P.D. Magnus's original version of this book went the other way.

and the following as their basic negation elimination rule:

$$
\begin{array}{lll}
m & & \neg\mathcal{A} \\
& & \\
n-1 & & \mathcal{B} \\
& & \\
n & & \neg\mathcal{B} \\
& & \\
& \mathcal{A} & \qquad \neg\text{E}^* \; m\text{--}n
\end{array}
$$

Using these two rules, we could have derived all of the rules governing negation and contradiction that we have taken as basic (i.e. $\perp$I, $\perp$E, $\neg$I 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?[2]

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 *tertium non datur* (i.e. TND) and *ex falso quodlibet* (i.e. $\perp$E). By treating these as separate rules in the proof system, we will be in a better position 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.

B.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 $\forall\chi\neg\mathcal{A}$ to $\neg\exists\chi\mathcal{A}$ and vice versa.[3]

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 $\mathcal{A}$ contains the name c, and contains no instances of the variable χ, and that we want to do the following:

$$
\begin{array}{ll}
m & \mathcal{A}(\ldots c \ldots c \ldots) \\
& \\
k & \exists\chi\mathcal{A}(\ldots \chi \ldots c \ldots)
\end{array}
$$

This is not yet permitted, since in this new system, we do not have the $\exists$I rule. We can, however, offer the following:

---

[2] Again, P.D. Magnus's original version of this book went the other way.
[3] Warren Goldfarb follows this line in *Deductive Logic*, 2003, Hackett Publishing Co.

m	$\mathcal{A}(\ldots c \ldots c \ldots)$	
$m+1$	$\quad \neg \exists \chi \mathcal{A}(\ldots \chi \ldots c \ldots)$	
$m+2$	$\quad \forall \chi \neg \mathcal{A}(\ldots \chi \ldots c \ldots)$	CQ $m+1$
$m+3$	$\quad \neg \mathcal{A}(\ldots c \ldots c \ldots)$	$\forall$E $m+2$
$m+4$	$\quad \bot$	$\bot$I m, $m+3$
$m+5$	$\neg\neg\exists \chi \mathcal{A}(\ldots \chi \ldots c \ldots)$	$\neg$I $m+1$–$m+4$
$m+6$	$\exists \chi \mathcal{A}(\ldots \chi \ldots c \ldots)$	DNE $m+5$

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:

m	$\exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$	
i	$\quad \mathcal{A}(\ldots c \ldots c \ldots)$	
j	$\quad \mathcal{B}$	
k	$\mathcal{B}$	

where c does not occur in any undischarged assumptions, or in $\mathcal{B}$, or in $\exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$. 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	$\exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$	
i	$\quad \mathcal{A}(\ldots c \ldots c \ldots)$	
j	$\quad \mathcal{B}$	
k	$\mathcal{A}(\ldots c \ldots c \ldots) \to \mathcal{B}$	$\to$I i–j
$k+1$	$\quad \neg \mathcal{B}$	
$k+2$	$\quad \neg \mathcal{A}(\ldots c \ldots c \ldots)$	MT k, $k+1$
$k+3$	$\quad \forall \chi \neg \mathcal{A}(\ldots \chi \ldots \chi \ldots)$	$\forall$I $k+2$
$k+4$	$\quad \neg \exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$	CQ $k+3$
$k+5$	$\quad \bot$	$\bot$I m, $k+4$
$k+6$	$\neg\neg\mathcal{B}$	$\neg$I $k+1$–$k+5$
$k+7$	$\mathcal{B}$	DNE $k+6$

We are permitted to use $\forall$I on line $k + 3$ because c does not occur in any undischarged assumptions or in $\mathcal{B}$. The entries on lines $k+4$ and $k+1$ contradict each other, because c does not occur in $\exists \chi \mathcal{A}(\ldots \chi \ldots \chi \ldots)$.

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 $\neg\forall \chi \mathcal{A}$ to $\exists \chi \neg \mathcal{A}$.)

Quick reference $\qquad$ C

C.1 Characteristic Truth Tables

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

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

C.2 symbolization

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	$\exists x(Fx\wedge\forall y(Fy\rightarrow x=y)\wedge\neg Gx)$

C.3   Using identity to symbolize quantities

There are at least _____ Fs.

one: $\exists x F x$

two: $\exists x_1 \exists x_2 (F x_1 \wedge F x_2 \wedge \neg x_1 = x_2)$

three: $\exists x_1 \exists x_2 \exists x_3 (F x_1 \wedge F x_2 \wedge F x_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3)$

four: $\exists x_1 \exists x_2 \exists x_3 \exists x_4 (F x_1 \wedge F x_2 \wedge F x_3 \wedge F x_4 \wedge$
$\neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_1 = x_4 \wedge \neg x_2 = x_3 \wedge \neg x_2 = x_4 \wedge \neg x_3 = x_4)$

n: $\exists x_1 \ldots \exists x_n (F x_1 \wedge \ldots \wedge F x_n \wedge \neg x_1 = x_2 \wedge \ldots \wedge \neg x_{n-1} = x_n)$

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 symbolization for 'there are at least $n + 1$ Fs'. Equivalently, we can offer:

one: $\forall x_1 \forall x_2 \big[(F x_1 \wedge F x_2) \rightarrow x_1 = x_2 \big]$

two: $\forall x_1 \forall x_2 \forall x_3 \big[(F x_1 \wedge F x_2 \wedge F x_3) \rightarrow (x_1 = x_2 \vee x_1 = x_3 \vee x_2 = x_3) \big]$

three: $\forall x_1 \forall x_2 \forall x_3 \forall x_4 \big[(F x_1 \wedge F x_2 \wedge F x_3 \wedge F x_4) \rightarrow$
$(x_1 = x_2 \vee x_1 = x_3 \vee x_1 = x_4 \vee x_2 = x_3 \vee x_2 = x_4 \vee x_3 = x_4) \big]$

n: $\forall x_1 \ldots \forall x_{n+1} \big[(F x_1 \wedge \ldots \wedge F x_{n+1}) \rightarrow (x_1 = x_2 \vee \ldots \vee x_n = x_{n+1}) \big]$

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: $\forall x \neg F x$

one: $\exists x \big[F x \wedge \forall y (F y \rightarrow x = y) \big]$

two: $\exists x_1 \exists x_2 \big[F x_1 \wedge F x_2 \wedge \neg x_1 = x_2 \wedge \forall y \big(F y \rightarrow (y = x_1 \vee y = x_2) \big) \big]$

three: $\exists x_1 \exists x_2 \exists x_3 \big[F x_1 \wedge F x_2 \wedge F x_3 \wedge \neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \neg x_2 = x_3 \wedge$
$\forall y \big(F y \rightarrow (y = x_1 \vee y = x_2 \vee y = x_3) \big) \big]$

n: $\exists x_1 \ldots \exists x_n \big[F x_1 \wedge \ldots \wedge F x_n \wedge \neg x_1 = x_2 \wedge \ldots \wedge \neg x_{n-1} = x_n \wedge$
$\forall y \big(F y \rightarrow (y = x_1 \vee \ldots \vee y = x_n) \big) \big]$

C.4 Basic deduction rules for TFL

Conjunction

$$
\begin{array}{l|l}
m & \mathcal{A} \\
n & \mathcal{B} \\
& \mathcal{A} \wedge \mathcal{B} \qquad \wedge\text{I } m, n
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \wedge \mathcal{B} \\
& \mathcal{A} \qquad \wedge\text{E } m
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \wedge \mathcal{B} \\
& \mathcal{B} \qquad \wedge\text{E } m
\end{array}
$$

Conditional

$$
\begin{array}{l|ll}
i & & \mathcal{A} \\
j & & \mathcal{B} \\
& \mathcal{A} \rightarrow \mathcal{B} & \rightarrow\text{I } i\text{--}j
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \rightarrow \mathcal{B} \\
n & \mathcal{A} \\
& \mathcal{B} \qquad \rightarrow\text{E } m, n
\end{array}
$$

Contradiction

$$
\begin{array}{l|l}
m & \mathcal{A} \\
n & \neg\mathcal{A} \\
& \bot \qquad \bot\text{I } m, n
\end{array}
$$

$$
\begin{array}{l|l}
m & \bot \\
& \mathcal{A} \qquad \bot\text{E } m
\end{array}
$$

Negation

$$
\begin{array}{l|ll}
i & & \mathcal{A} \\
j & & \bot \\
& \neg\mathcal{A} & \neg\text{I } i\text{--}j
\end{array}
$$

Tertium non datur

$$
\begin{array}{l|ll}
i & & \mathcal{A} \\
j & & \mathcal{B} \\
k & & \neg\mathcal{A} \\
l & & \mathcal{B} \\
& \mathcal{B} & \text{TND } i\text{--}j,\ k\text{--}l
\end{array}
$$

Disjunction

$$
\begin{array}{l|l}
m & \mathcal{A} \\
& \mathcal{A} \vee \mathcal{B} \qquad \vee\text{I } m
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \\
& \mathcal{B} \vee \mathcal{A} \qquad \vee\text{I } m
\end{array}
$$

$$
\begin{array}{l|ll}
m & \mathcal{A} \vee \mathcal{B} \\
i & & \mathcal{A} \\
j & & \mathcal{C} \\
k & & \mathcal{B} \\
l & & \mathcal{C} \\
& \mathcal{C} & \vee\text{E } m,\ i\text{--}j,\ k\text{--}l
\end{array}
$$

Biconditional

$$
\begin{array}{l|l}
i & \quad \mathcal{A} \\
  & \quad \overline{\phantom{\mathcal{A}}} \\
j & \quad \mathcal{B} \\[4pt]
k & \quad \mathcal{B} \\
  & \quad \overline{\phantom{\mathcal{B}}} \\
l & \quad \mathcal{A} \\[4pt]
  & \mathcal{A} \leftrightarrow \mathcal{B} \qquad \leftrightarrow\text{I } i\text{--}j,\, k\text{--}l
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \leftrightarrow \mathcal{B} \\
n & \mathcal{A} \\
  & \mathcal{B} \qquad\qquad \leftrightarrow\text{E } m,\, n
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \leftrightarrow \mathcal{B} \\
n & \mathcal{B} \\
  & \mathcal{A} \qquad\qquad \leftrightarrow\text{E } m,\, n
\end{array}
$$

C.5 Derived rules for TFL

Disjunctive syllogism

$$
\begin{array}{l|l}
m & \mathcal{A} \vee \mathcal{B} \\
n & \neg\mathcal{A} \\
  & \mathcal{B} \qquad\quad \text{DS } m,\, n
\end{array}
$$

$$
\begin{array}{l|l}
m & \mathcal{A} \vee \mathcal{B} \\
n & \neg\mathcal{B} \\
  & \mathcal{A} \qquad\quad \text{DS } m,\, n
\end{array}
$$

Reiteration

$$
\begin{array}{l|l}
m & \mathcal{A} \\
  & \mathcal{A} \qquad \text{R } m
\end{array}
$$

Modus Tollens

$$
\begin{array}{l|l}
m & \mathcal{A} \rightarrow \mathcal{B} \\
n & \neg\mathcal{B} \\
  & \neg\mathcal{A} \qquad\quad \text{MT } m,\, n
\end{array}
$$

Double-negation elimination

$$
\begin{array}{l|l}
m & \neg\neg\mathcal{A} \\
  & \mathcal{A} \qquad\quad \text{DNE } m
\end{array}
$$

De Morgan Rules

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

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

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

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

C.6 Basic deduction rules for FOL

Universal elimination

$$
\begin{array}{l|l}
m & \forall x \mathcal{A}(\ldots x \ldots x \ldots) \\
& \mathcal{A}(\ldots c \ldots c \ldots) \qquad \forall \text{E } m
\end{array}
$$

Universal introduction

$$
\begin{array}{l|l}
m & \mathcal{A}(\ldots c \ldots c \ldots) \\
& \forall x \mathcal{A}(\ldots x \ldots x \ldots) \qquad \forall \text{I } m
\end{array}
$$

c must not occur in any undischarged assumption
x must not occur in $\mathcal{A}(\ldots c \ldots c \ldots)$

Identity introduction

$$
\begin{array}{l|l}
& c = c \qquad =\text{I}
\end{array}
$$

Existential introduction

$$
\begin{array}{l|l}
m & \mathcal{A}(\ldots c \ldots c \ldots) \\
& \exists x \mathcal{A}(\ldots x \ldots c \ldots) \qquad \exists \text{I } m
\end{array}
$$
x must not occur in $\mathcal{A}(\ldots c \ldots c \ldots)$

Existential elimination

$$
\begin{array}{l|l}
m & \exists x \mathcal{A}(\ldots x \ldots x \ldots) \\
i & \quad \begin{array}{|l} \mathcal{A}(\ldots c \ldots c \ldots) \\ \hline \end{array} \\
j & \quad \begin{array}{|l} \mathcal{B} \end{array} \\
& \mathcal{B} \qquad \exists \text{E } m, i\text{–}j
\end{array}
$$

c must not occur in any undischarged assumption, in $\exists x \mathcal{A}(\ldots x \ldots x \ldots)$, or in $\mathcal{B}$

Identity elimination

$$
\begin{array}{l|l}
m & a = b \\
n & \mathcal{A}(\ldots a \ldots a \ldots) \\
& \mathcal{A}(\ldots b \ldots a \ldots) \qquad =\text{E } m, n
\end{array}
\qquad
\begin{array}{l|l}
m & a = b \\
n & \mathcal{A}(\ldots b \ldots b \ldots) \\
& \mathcal{A}(\ldots a \ldots b \ldots) \qquad =\text{E } m, n
\end{array}
$$

C.7 Derived rules for FOL

$$
\begin{array}{l|l}
m & \forall x \neg \mathcal{A} \\
& \neg \exists x \mathcal{A} \qquad \text{CQ } m
\end{array}
$$

$$
\begin{array}{l|l}
m & \neg \exists x \mathcal{A} \\
& \forall x \neg \mathcal{A} \qquad \text{CQ } m
\end{array}
$$

$$
\begin{array}{l|l}
m & \exists x \neg \mathcal{A} \\
& \neg \forall x \mathcal{A} \qquad \text{CQ } m
\end{array}
$$

$$
\begin{array}{l|l}
m & \neg \forall x \mathcal{A} \\
& \exists x \neg \mathcal{A} \qquad \text{CQ } m
\end{array}
$$

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 University Lecturer, and Fellow of St John's College, at the University of Cambridge. His first book, *The Limits of Realism*, was published by Oxford University Press in 2013.

Kathryn Lindeman is an assistant professor of philosophy at Saint Louis University. Her primary research is in ethics, metaethics, and epistemology.