Notes to readers

- What follows is a (revised) draft of the opening chapters for a planned second edition of my *Introduction to Formal Logic* (CUP, 2003). This is a textbook aimed at first-year philosophy students: if all goes well, the second edition will appear in 2019 in the *Cambridge Introductions to Philosophy* series. These preliminary chapters aim to introducing some key concepts in a relaxed way, and should be useful background whatever formal logic course you then go on to take. The chapters go slowly, you might think *very* slowly, particularly at the outset!
- Work no doubt still needs to be done improving/extending the exercise sets at the end of the chapters. There will eventually be answers to the exercises on the web.
- At this stage, all kinds of comments (other than ones that mean, in effect, 'You are writing the wrong book!') are most welcome. Obviously I want to hear about any typos you spot. But in addition, it is very useful to hear e.g. about passages that you found obscure or more difficult than others, and passages you thought were possibly misleading. And if you are not a native English-speaker, do note any words or turns of phrase that you found puzzling.
- Spelling follows British English rather than North American English conventions. So I write e.g. 'fulfil' rather than 'fulfill', 'skilful' rather than 'skillful', etc. I aim to systematically use 'ize' endings where appropriate. (But still, if you think what I've written is a typo, do say so – I'd rather you over-corrected than under-corrected!)
- One issue about punctuation. Suppose we have some introductory words followed by some indented displayed material. Should the introductory words end with a colon or not? My general line is this: if the introduction is a complete clause, I use a colon; if the thought runs on straight into the indented material, I don't.
- I use 'they'/'them' as gender-neutral singular pronouns. (I don't entirely like this but prefer it to alternatives. Let me know of cases that really grate!)
- Unresolved references of the kind '**??**' are forward-looking to chapters or sections that aren't included in the current chapters and will be resolved in due course.
- Comments and corrections please to ps218 at cam dot ac dot uk – if you do send any, it could be *very* helpful if you mention the date of this draft.

An Introduction to Formal Logic

Second edition

Peter Smith

October 30, 2018

Contents

1 What is deductive logic?

The core business of logic is the *systematic evaluation of arguments for internal cogency*. And the kind of internal cogency that will especially concern us in this book is *logical validity*.

These brief headlines leave everything to be explained! What do we mean here by 'argument'? What do we mean by 'internal cogency'? What do we mean, more particularly, by 'logical validity'? And what kinds of 'systematic' evaluation of arguments are possible? This introductory chapter makes a gentle start on answering these questions.

1.1 What is an argument?

By 'argument' we mean a chain of reasoning, short or long, in support of some conclusion. So we must distinguish arguments from mere disagreements and disputes. The children who shout at each other 'You did', 'I didn't', 'Oh yes, you did', 'Oh no, I didn't', are certainly disagreeing: but they are not *arguing* in our sense – they are not yet giving any reasons in support of one claim or the other.

Reason-giving arguments are the very stuff of all serious enquiry, whether it is philosophy or physics, economics or experimental psychology. But of course, episodes of reasoning equally feature in everyday, street-level, enquiry into the likely winner of next month's election or the best place to train as a lawyer, or into what explains our team's losing streak. We want our opinions to be true; which means that we should aim to have good reasons backing up our opinions, so raising the chances of getting things right. That in turn means that we have an interest in being skilful reasoners, using arguments which really do support their conclusions.

1.2 Kinds of evaluation

Logic, then, is concerned with evaluating stretches of reasoning. Take a really simple example. Suppose you hold

A (1) All philosophers are eccentric.

I then introduce you to Jack, telling you that he is a philosopher. So you come to believe

(2) Jack is a philosopher.

Putting these two thoughts together, you infer

(3) Jack is eccentric.

And the first point to make is that this little bit of reasoning, argument **A**, can now be evaluated along two quite independent dimensions:

First, we can ask whether **A**'s *premisses* (1) and (2) are true. In other words, are the 'inputs' to your inference correct? (1) is in fact very disputable. And perhaps I have made a mistake, and (2) is false as well.

Second, we can ask about the quality of the *inference step*, the move which takes us from the premisses (1) and (2) to the *conclusion* (3). In this particular case, the inference step is surely absolutely compelling: the conclusion really does follow from the premisses. We have agreed that it may be open to question whether (1) and (2) are actually true. However, if they *are* assumed to be true (assumed 'for the sake of argument', as we say), then we have to agree that (3) is true too. There's just no way that (1) and (2) could be true and yet (3) false. To assert that Jack is a philosopher and that all philosophers are eccentric, but go on to deny that Jack is eccentric, would be implicitly to contradict yourself.

Generalizing, it is one thing to consider whether an argument starts from true premisses; it is another thing to consider whether it moves on by reliable inference steps. Yes, we typically want our arguments to pass muster on both counts. We typically want *both* to start from true premisses *and* to reason by steps which can be relied on to take us to further truths. But it is important to emphasize that these are distinct aims.

The premisses of arguments can be about all sorts of topics: their truth is usually no business of the logician. If we are arguing about historical matters, then it is the historian who is the expert about the truth of our premisses; if we are arguing about some question in physics, then the physicist is the one who might know whether our premisses are true; and so on. The central concern of logic, by contrast, is not the truth of initial premisses but the way we argue from a given starting point. It is in this sense that logic is concerned with the 'internal cogency' of our reasoning. The logician's key question is whether an argument's premisses, supposing that we accept them, really do give us good grounds for accepting its conclusion.

1.3 Deduction vs. induction

(a) If **A**'s premisses are true, then its conclusion is guaranteed to be true too. Here is a similar case:

B (1) Either Jill is in the library or she is in the bookshop.
(2) Jill isn't in the library.
So (3) Jill is in the bookshop.

Who knows whether the initial assumptions, the two premisses, are true or not? But we can immediately see that the inference step is again completely watertight. If premisses B(1) and B(2) are both true, then B(3) cannot conceivably fail to be true.

Now consider the following contrasting case. Here you are, sitting in your favourite café. Unrealistic philosophical scepticism apart, you are thoroughly confident that the cup of coffee you are drinking is not going to kill you – for if you weren't really confident, you wouldn't be calmly sipping as you read this, would you? What justifies your confidence?

Well, you believe the likes of:

C (1) Cups of coffee from GreatBeanz that looked and tasted just fine haven't killed anyone in the past.
(2) This present cup of GreatBeanz coffee looks and tastes just fine.

These premisses, or something like them, sustain your cheerful belief that

(3) This present cup of GreatBeanz coffee won't kill you.

The inference that moves from the premisses C(1) and C(2) to the conclusion C(3) is, in normal circumstances, surely perfectly reasonable: other things being equal, the facts recorded in C(1) and C(2) do give you excellent grounds for believing that C(3) is true. However – and here is the quite crucial contrast with the earlier 'Jack' and 'Jill' examples – it is not the case that the truth of C(1) and C(2) absolutely guarantees C(3) to be true too.

Perhaps someone has slipped a slow-acting tasteless poison into the coffee, just to make the logical point that facts about how things have always been in the past don't guarantee that the trend will continue in the future.

Fortunately for you, C(3) is no doubt true. The tasteless poison is a fantasy. Still, it is a *coherent* fantasy. It illustrates the point that your grounds C(1) and C(2) for the conclusion that the coffee is safe to drink are strictly speaking quite compatible with the falsity of that conclusion. Someone who agrees to C(1) and C(2) and yet goes on to assert the opposite of C(3) might be saying something highly improbable, but they won't actually be contradicting themselves. We can make sense of the idea of C(1) and C(2) being true and yet C(3) false.

In summary then, there is a fundamental difference between the 'Jack' and 'Jill' examples on the one hand, and the 'coffee' example on the other. In the 'Jack' and 'Jill' cases, the premisses absolutely guarantee the conclusion. There is no conceivable way that A(1) and A(2) could be true and yet A(3) false: likewise, if B(1) and B(2) are true then B(3) just has to be true too. Not so with the 'coffee' case: it is conceivable that C(1) and C(2) are true while C(3) is false. What has happened in the past is a very good guide to what will happen next (and what else can we rely on?): but reasoning from past to future isn't completely watertight.

(b) We need some terminology to mark this fundamental difference. We will introduce it informally for the moment:

> If an inference step from premisses to a conclusion is completely watertight, i.e. if the truth of the premisses absolutely guarantees the truth of the conclusion, then we say that it is *deductively valid*.
>
> Equivalently, when an inference step is deductively valid, we will say that its premisses *deductively entail* its conclusion.

Hence the inferential moves in **A** and **B** count as being deductively valid. Contrast the coffee argument **C**. That argument involves reasoning from past cases to a new case in a way which leaves room for error, however unlikely. This kind of extrapolation from the past to the future, or more generally from some sample cases to further cases, is standardly called *inductive*. The inference in **C** might be inductively strong – meaning that the conclusion is highly probable, assuming the premisses are true – but the inference is not deductively valid.

We should stress then that the deductive/inductive distinction is *not* the distinction between good and bad reasoning. The 'coffee' argument is a perfectly decent one. It involves the sort of usually reliable reasoning to which we trust our lives, day in, day out (what else can we do?). It is just that the inference step here doesn't completely guarantee that the conclusion is true, even assuming that the stated premisses are true.

(c) What makes for reliable (or reliable enough) inductive inferences is a very important and decidedly difficult topic. But it is not our topic in this book, which is deductive logic. That is to say, we will be here concentrating on the assessment of arguments which aim to use deductively valid inferences, where the premisses *are* supposed to strictly entail the conclusion.

We will give a sharper definition of the general notion of deductive validity at the beginning of the next chapter, §2.1. Later, by the time we get to §6.2, we will have the materials to hand to define a rather narrower notion, which following tradition we will call *logical* validity. And this narrower notion of logical validity will then become our main focus in the remainder of the book. But for the next few chapters, we continue to work with the wider initial notion that we've called deductive validity.

1.4 Just a few more examples

The 'Jack' and 'Jill' arguments are examples where the inference steps are obviously deductively valid. Compare this next argument:

D (1) All Republican voters support capital punishment.
(2) Jo supports capital punishment.
So (3) Jo is a Republican voter.

The inference step here is equally obviously invalid. Even if D(1) and D(2) are true, D(3) doesn't follow. Maybe lots of people in addition to Republican voters support capital punishment, and Jo is one of them.

How about the following argument?

E (1) Most Irish are Catholics.
(2) Most Catholics oppose abortion on demand.
So (3) At least some Irish oppose abortion on demand.

Leave aside the question of whether the premisses are in fact correct (that's not a matter for logicians: it needs sociological investigation to determine the distribution of religious affiliation among the Irish, and to find out what proportion of Catholics support their church's official teaching about abortion). What we can ask here – from our armchairs, so to speak – is whether the inference step is deductively valid: if the premisses are true, then must the conclusion be true too?

Well, whatever the facts of the case, it is at least conceivable that the Irish are a tiny minority of the Catholics in the world. And it could also be that nearly all the other (non-Irish) Catholics oppose abortion, and hence most Catholics do, even though *none* of the Irish oppose abortion. But then E(1) and E(2) would be true, yet E(3) false. So the truth of the premisses doesn't by itself absolutely guarantee the truth of the conclusion (there are possible situations in which the premisses would be true and the conclusion false). The inference step here can't be deductively valid.

Here's another trite argument: is the inference step deductively valid this time?

F (1) Some philosophy students admire all logicians.
(2) No philosophy student admires any rotten lecturer.
So (3) No logician is a rotten lecturer.

With a little thought you should arrive at the right answer here too (we will return to this example in Chapter 3).

Still, at the moment, faced with examples like our last three, all you can do is to cast around hopefully, trying to work out somehow or other whether the truth of the premisses *does* guarantee the truth of the conclusion. It would obviously be good to be able to proceed more systematically and to have some *general* techniques for evaluating arguments for deductive validity. That's what logical theory aims to provide.

Indeed, ideally, we would like techniques that work mechanically, that can be applied to settle questions of validity as routinely as we can settle simple arithmetical questions by calculation. We will have to wait to see how far this is possible. For the moment, we will just say a little more about what makes any kind of more systematic approach possible (whether mechanical or not).

1.5 Generalizing

(a) Here again is our first sample mini-argument with its obviously valid inference step:

A (1) All philosophers are eccentric.
(2) Jack is a philosopher.
So (3) Jack is eccentric.

Now compare it with the following arguments:

A′ (1) All logicians are cool.
(2) Russell is a logician.
So (3) Russell is cool.

A″ (1) All robins have red breasts.
(2) Tweety is a robin.
So (3) Tweety has a red breast.

A‴ (1) All post-modernists write nonsense.
(2) Derrida is a post-modernist.
So (3) Derrida writes nonsense.

We can keep going on and on, churning out arguments to the same pattern, all involving equally valid inference steps.

It is plainly no accident that these arguments share the property of being internally cogent. Comparing these examples makes it clear that the deductive validity of the inference step in the original argument **A** hasn't anything specifically to do with philosophers or with the notion of being eccentric. Likewise the validity of the inference step in argument **A′** hasn't anything specifically to do with logicians. And so on. There is a general principle involved here, which we could put like this:

> From a pair of premisses, one saying that all things of a certain kind have a given property, the other saying that a particular individual is of the former kind, we can validly infer the conclusion that the individual in question has the latter property.

(b) However, that wordy version is not the most perspicuous way of representing the shared inferential principle at work in the **A**-family. Focusing on arguments couched in English, we can instead say something like this:

Any inference step of the form

All F are G
n is F
So: n is G

is deductively valid.

Here the italic letters '*n*', '*F*', '*G*' are being used to help exhibit a skeletal pattern of argument. We can think of '*n*' as standing in for a name for some person or thing. While '*F*' and '*G*' stand in for expressions which pick out kinds of things (like 'philosopher' or 'robin') or properties (like 'eccentric'). A bit more loosely, we can also use 'is *F*' to stand in for other expressions attributing properties like 'has a red breast' or 'writes nonsense'. It then doesn't matter how we flesh out the italicized argument template or *schema*, as we might call it. Any sensible enough way of substituting suitable expressions for '*n*', '(is) *F*' and '(is) *G*', and then tidying the result into reasonable English, will yield another argument with a deductively valid inference step.

Some forms of inference – like the one just illustrated – are deductively reliable, meaning that every inference step which is an instance of the displayed schematic pattern is valid. Other forms aren't reliable. Consider the following pattern of inference:

Most F are G
Most G are H
So: At least some F are H.

This is the type of inference involved in the 'Irish' argument **E**, and we now know that it isn't trustworthy.

(c) We said at the outset that logic aims to be a systematic study. We can now begin to see how to get some generality into the story. Noting that the same form of inference step can feature in many different particular arguments, we can aim to examine generally reliable forms of inference. And we can then explore how such reliable forms relate to each other and to the particular arguments which exploit them.

So there will be a lot more in later chapters on this idea of general principles of inference, principles which we can display using schematic patterns. But first, a quick summary.

1.6 Summary

We can evaluate a piece of reasoning along two different dimensions. We can ask whether its premisses are actually true. And we can ask whether the inference from the premisses is cogent – i.e., assuming the premisses are true, do they really support the truth of the conclusion? Logic is concerned with the second dimension of evaluation.

We are setting aside inductive arguments (and other kinds of non-conclusive reasoning). We will be concentrating on arguments involving inference steps that purport to be deductively valid. In other words, we are going to be concerned with deductive logic, the study of inferences that aim to strictly guarantee their conclusions, assuming the truth of their premisses.

Arguments typically come in families whose members share good or bad types of inferential move; by looking at such general patterns or forms of inference we can hope to make logic more systematic.

Exercises 1

We return to the key idea of *forms* of inference in Chapter 3. For now, a handful of examples about elementary valid/invalid inferences.

Note, by ‘conclusion’ we do not mean what concludes a passage of reasoning in the sense of what is stated at the end. We mean what the reasoning aims to establish – and that might in fact be stated at the outset. Likewise, ‘premiss’ does not mean (contrary to what the Concise Oxford Dictionary says!) ‘a previous statement from which another is inferred’. Reasons supporting a certain conclusion, i.e. the inputs to an inference, might well be given *after* that target conclusion has been stated. And the move from supporting reasons to conclusions can be signalled by *inference markers* other than ‘So’.

Indicate the premisses, inference markers, and conclusions of the following one-step arguments. Which of these arguments do you suppose involve deductively valid reasoning? Why? (Just improvise, and answer the best you can!)

(1) Whoever works hard at logic does well. Accordingly, if Russell works hard at logic, he does well.

(2) Most politicians are corrupt. After all, most ordinary people are corrupt – and politicians are ordinary people.

(3) It will snow tonight – because the snow clouds show up clearly on the weather satellite, heading this way.

(4) Anyone who is well prepared for the exam, even if she doesn’t get an A grade, will at least get a B. Jane is well prepared, so she will get at least a B grade.

(5) John is taller than Mary, and Jane is shorter than Mary. So John is taller than Jane.

(6) At eleven, Fred is always either in the library or in the coffee bar. And assuming he’s in the coffee bar, he’s drinking an espresso. Fred was not in the library when I looked at eleven. So he was drinking an espresso then.

(7) The Democrats will win the election. For the polls put them 20 points ahead, and no party can ever overturn even a lead of 10 points with only a week to go to polling day.

(8) Dogs have four legs. Fido is a dog. Therefore Fido has four legs.

(9) Jekyll isn't the same person as Hyde. The reason is that no murderers are sane – but Hyde is a murderer, and Jekyll is certainly sane.

(10) All the slithy toves did gyre and gimble in the wabe. Some mome raths are slithy toves. Hence some mome raths did gyre and gimble in the wabe.

(11) No experienced person is incompetent. Jenkins is always blundering. No competent person is always blundering. Therefore Jenkins is inexperienced.

(12) Many politicians take bribes. Most politicians have extra-marital affairs. So many people who take bribes have extra-marital affairs.

(13) Some but not all philosophers are logicians. All logicians are clever. Hence some but not all philosophers are clever.

(14) Kermit is green all over. Hence Kermit is not red all over.

(15) Miracles cannot happen. Since, by definition, a miracle is an event incompatible with the laws of nature. And everything that happens must be consistent with the laws of nature.

(16) (Lewis Carroll) Babies cannot manage crocodiles. For babies are illogical – but illogical persons are despised, and nobody is despised who can manage a crocodile.

(17) (Lewis Carroll again) No interesting poems are unpopular among people of real taste. No modern poetry is free from affectation. All your poems are on the subject of soap bubbles. No affected poetry is popular among people of real taste. Only a modern poem would be on the subject of soap bubbles. Therefore all your poems are uninteresting.

(18) [For those who know the tiniest amount of maths!] Define $P(k)$ to mean that the sum of the first k positive integers is $k(k + 1)$. Then $P(0)$. And for any positive integer n, if $P(n)$ then $P(n + 1)$. Hence, for any positive integer n, $P(n)$.

(19) 'If we found by chance a watch or other piece of intricate mechanism we should infer that it had been made by someone. But all around us we do find intricate pieces of natural mechanism, and the processes of the universe are seen to move together in complex relations; we should therefore infer that these too have a maker.'

(20) 'I can doubt that the physical world exists. I can even doubt whether my body really exists. I cannot doubt that I myself exist. So I am not my body.'

2 Validity and soundness

Our first chapter introduced the idea of an inference step being deductively valid – or, in equivalent terminology, the idea of some premisses entailing a conclusion. This chapter explores this notion of validity/entailment a bit further, though still in an informal way. We also emphasize the special centrality of deductive reasoning in serious enquiry.

2.1 Validity defined again

(a) We said in §1.3(b) that an inference step is deductively valid if it is completely watertight – in other words, assuming that the inference's premisses are true, its conclusion is absolutely guaranteed to be true as well. But talk of an inference being 'watertight' (and likewise talk of a conclusion being 'guaranteed' to be true) is too metaphorical for comfort. Let's try to pin down our intended idea of validity rather more carefully, even if still somewhat informally.

So here is a less metaphorical definition, already hinted at:

> An inference step is *deductively valid* if and only if there is no possible situation in which its premisses would be true and its conclusion false.
>
> Equivalently, in such a case, we will say the inference's premisses (deductively) *entail* its conclusion.

We can take the plural 'premisses', here and in similar statements, to cover the one-premiss case too. This pretty standard definition characterizes what is often called the *classical* concept of validity. It is, however, only as clear as the notion of a 'possible situation'. So we certainly need to pause over this.

(b) Consider the following bit of reasoning:

A Jo jumped out of a twentieth floor window (without parachute, safety net, etc.) and fell unimpeded onto a concrete pavement. So she was injured.

Let's grant that, with the laws of nature as they are, there is no way in which the premiss could be true and the conclusion false (assume that we are talking about here on Earth, that Jo is an adult human, not a beetle, and so on). In this world as it actually is, falling unimpeded onto concrete from twenty floors up will always produce serious – very probably fatal – injury. Does that make the inference from **A**'s premiss to its conclusion deductively valid?

No. It isn't, let's agree, *physically* possible in our actual circumstances to jump without being injured. But we can coherently, without self-contradiction, conceive of a situation in which the laws of nature are different or are miraculously suspended, and someone jumping from twenty floors up will float delicately down like a feather. We can imagine a capricious deity bringing about such a situation. In this very weak sense, then, it *is* possible that **A**'s premiss should be true while its conclusion is false. And that is enough for the inference to be deductively invalid.

(c) So our definition of validity is to be read like this: a deductively valid inference is one where it is not possible even in the most generously inclusive sense – it is not even coherently conceivable – that the inference's premisses are true and its conclusion false.

Let's elaborate a little more. We ordinarily talk of many different kinds of possibility – about what is physically possible (given the laws of nature), what is technologically possible (given our current engineering skills), what is politically possible (given the make-up of the electorate), what is financially possible (given the state of your bank balance), and so on. These different kinds of possibility are related in various ways. For example, something that is technologically possible has to be physically possible, but may not be financially possible for you (or indeed anyone else) to engineer.

But now note that if some situation is physically possible (or technologically possible, or politically possible, etc.) then, at the very least, a description of that situation has to be internally coherent. In other words, if some situation is M-ly possible (for any suitable modifier 'M'), then it must also be possible in the thin sense that the very idea of such a situation obtaining isn't ruled out as nonsense. The notion of possibility we need in defining deductive validity is possibility in this weakest, most inclusive, sense. And from now on, when we talk in an unqualified way about possibility, it is this very weak sense that we will have in mind.

(d) Kinds of possibility go together with matching kinds of necessity – 'necessary' means 'not-possibly-not'. For example, it is politically necessary to implement the result of a referendum if and only if it is not politically possible *not* to implement that result. It is legally necessary for an enforceable will to be signed if and only if it is not legally possible for a will to be enforceable and *not* signed. It is physically necessary that massive objects gravationally attract each other if and only if it is not physically possible that those objects *not* attract each other. To generalize: we can massage such equivalences into the standard form

It is M-ly necessary that C if and only if it is not M-ly possible that *not*-C,

where 'C' stands in for some statement or proposition, and '*not*-C' stands in for the *denial* of that proposition (so asserting *not*-C is equivalent to saying it is false that C).

In the same way, our very weak inclusive notion of possibility goes with a correspondingly very strong notion of necessity: it is necessary that C in this sense just in case it is not even weakly possible that *not*-C. Putting it another way,

> It is necessarily true that C if and only if it is true that C in every possible situation, in our most inclusive sense of 'possible'.

And from now on, when we talk in an unqualified way about necessity, it is this very strong sense that we will have in mind.

And yes, being necessarily true is a *very* strong requirement, as strong as can be. However, it is one that can be met in some cases. For example, it is necessary in this sense that Jill is not both married and unmarried (for in no coherently conceivable situation is Jill both married and not married). Again, it is necessary in our strong sense that all triangles have three sides (in any coherently conceivable situation, the triangular figures are three-sided). Likewise, it is necessary that no bachelor is married. And it is necessary that whenever a mother smiles, a parent smiles.

It is equally necessary that if all philosophers are eccentric and Jack is a philosopher, then Jack is eccentric. And this last example illustrates a general point, which gives us an equivalent definition of validity:

> An inference step is deductively valid if and only if it is necessary, in the strongest sense, that if the inference's premisses are true, so is its conclusion. Or putting it another way, deductively valid inferences are *necessarily truth-preserving*.

2.2 Consistency, validity, and equivalence

We need a term for the kind of thing that can feature as a premiss or conclusion in an argument – let's use the traditional '*proposition*'.

We will return in §§7.4–7.6 to consider the tricky question of the nature of propositions: for the moment, we just assume that propositions can sensibly be said to be true or false in various situations (unlike commands, questions, etc.). That's all we need, if we are to introduce two more notions that swim in the same conceptual stream as the notion of validity – two notions which again concern the truth of propositions across different possible situations.

(a) First, we will say:

> Some propositions are jointly *consistent* with each other if and only if there is a possible situation in which these propositions are all true together.

Or to put it another way, some propositions are jointly *inconsistent* if and only if there is no possible situation in which they are all true together. Both times, we mean 'possible' in the weakest, most inclusive, sense.

Suppose then that (1) an inference is valid, meaning that (2) there is no possible situation in which its premisses would be true and conclusion false. Given our new definition, (2) is equivalent to saying that (3) these premisses and the *denial* of the conclusion are jointly inconsistent. We can therefore offer an alternative definition of the classical conception of validity, worth highlighting:

> An inference step is deductively valid if and only if its premisses taken together with the denial of the conclusion are inconsistent.

This means that it matters little whether we develop logical theory by taking the notion of validity as basic or the notion of consistency.

(b) Here is the second new notion we want:

> Two propositions are *equivalent* if and only if they are true in exactly the same possible situations.

Again this is closely tied to the notion of a valid inference, as follows:

The propositions *A* and *B* are equivalent if and only if *A* entails *B* and *B* entails *A*.

Why so? Suppose *A* and *B* are equivalent. Then in any situation in which *A* is true, *B* is true – i.e. *A* entails *B*. Exactly similarly, *B* entails *A*. Conversely, suppose *A* entails and is entailed by *B*; then in any situation in which *A* is true, *B* is true and vice versa – i.e. *A* and *B* are equivalent.

2.3 Validity, truth, and the invalidity principle

(a) To assess whether an inference step is deductively valid, it is usually not enough to look at what is the case in the actual world; we also need to consider alternative possible situations, alternative ways things might conceivably have been – or, as some say, alternative possible worlds.

Take, for example, the following argument:

B (1) No Welshman is a great poet.
(2) Shakespeare is a Welshman.
So (3) Shakespeare is not a great poet.

The propositions here are all false about the actual world. But that doesn't settle whether the premisses of **B** entail its conclusion. In fact, the inference step here is a valid one; there is indeed no situation at all – even a merely possible one, in the weakest sense – in which the premisses would be true and the conclusion false. In other words, any possible situation which *did* make the premisses true (Shakespeare being brought up some miles to the west, and none of the Welsh, now including Shakespeare, being particularly good at verse), would also make the conclusion true.

Here are two more arguments stamped out from the same mould:

C (1) No human being is a dinosaur.
(2) Bill Clinton is a human being.
So (3) Bill Clinton is not a dinosaur.

D (1) No one whose middle name is 'William' is a Democrat.
(2) George W. Bush's middle name is 'William'.
So (3) George W. Bush is not a Democrat.

These two arguments involve deductively valid inferences of the same type – in case **C** taking us from true premisses to a true conclusion, in case **D** taking us, as it happens, from false premisses to a true conclusion.

You might wonder about the last case: how can a truth be validly inferred from two falsehoods? But note again: validity is only a matter of internal cogency. Someone who believes D(3) on the basis of the premisses D(1) and D(2) would be using a reliable form of inference, but they would be arriving at a true conclusion by sheer luck. Their claimed grounds D(1) and D(2) for believing D(3) would have nothing to do with why D(3) in fact happens to be true. (Aristotle noted the point. Discussing valid arguments with false premisses, he remarks that the conclusion may yet be correct, but "only in respect to the fact, not to the reason".)

There can be mixed cases too, where valid inferences have some true and some false premisses. So, allowing for these, we have in summary:

> A deductively valid inference can have actually true premisses and a true conclusion, (some or all) actually false premisses and a false conclusion, or (some or all) false premisses and a true conclusion.
>
> The only combination ruled out by the definition of deductive validity is a valid inference step's having all true premisses and a false conclusion. Deductive validity is about the necessary preservation of truth – so a deductively valid inference cannot take us from actually true premisses to an actually false conclusion.

That last point is worth really emphasizing. Here's an equivalent version:

> *The invalidity principle:* An inference step with actually true premisses and an actually false conclusion must be deductively invalid.

We will see in Chapter 5 how this principle gets to do important work when combined with the already-noted fact that arguments come in families sharing the same kind of inference step.

(b) So much for the deductively valid inference steps. What about the invalid ones? What combinations of true/false premisses and true/false conclusions can these have? A moment's reflection shows that

> A deductively invalid inference step can have any combination of true or false premisses and a true or false conclusion.

Take for example the silly inference

E (1) My mother's maiden name was 'Moretti'.
(2) Every computer I've ever bought is a Mac.
So (3) The Pavel Haas Quartet is my favourite string quartet.

Obviously, there is a conceivable situation in which the premisses are true and conclusion false. But what about the *actual* truth/falsity of the premisses and conclusion? I'm not going to tell you. But plainly, any combination of true or false premisses and a true or false conclusion here is compatible with the invalidity of that silly inference.

Hence: having true premisses and a false conclusion is *enough* to make an inference invalid – in other words, it is a sufficient condition for invalidity. But that isn't *required* for being invalid – it isn't a necessary condition.

2.4 Inferences and arguments

(a) So far, we have spoken of *inference steps in arguments* as being deductively valid or invalid. But it is more common simply to describe *arguments* as being valid or invalid. In this usage, we say that an argument is deductively valid if and only if the inference step from its premisses to its conclusion is valid. (Or at least, that's what we say about one-step arguments like the toy examples we've been looking at: we'll consider multi-inference arguments later.)

Talking of valid arguments in this way is absolutely standard. But it can mislead beginners. After all, saying that an argument is 'valid' can sound like an all-in endorsement. So let's stress the point: to say that a (one-step) argument is valid in our sense is *only* to commend the cogency of the inferential move between premisses and conclusion, is only to say that the conclusion really does follow from the premisses. A deductively valid argument, i.e. one that is internally cogent, can still have premisses that are quite hopelessly false.

For example, take this argument:

F (1) Whatever Donald Trump says is true.
(2) Donald Trump says that the Flying Spaghetti Monster exists.
So (3) It is true that the Flying Spaghetti Monster exists.

Since the *inference step* here is evidently valid, it follows that this *argument* counts as valid, in our standard usage. Which might very well strike the unwary as a distinctly odd thing to say!

You just have to learn to live with this way of speaking of valid arguments. But then we need, of course, to have a different term for arguments that *do* deserve all-in endorsement – arguments which both start from truths and proceed by deductively cogent inference steps. The usual term is 'sound'. So:

> A (one-step) argument is *deductively valid* if and only if the inference step from the premisses to the conclusion is valid.
>
> A (one-step) argument is (deductively) *sound* if and only if it has all true premisses and the inference step from those premisses to the conclusion is valid.

A few older texts annoyingly use 'sound' to mean what we mean by 'valid'. But everyone agrees there is a key distinction to be made between mere deductive cogency and the all-in virtue of making a cogent inference *and* having true premisses; there's just a divergence over how this agreed distinction should be labelled. Our usage of the labels, according to which **F** is valid but not sound, follows the modern convention.

(b) When we are deploying deductive arguments, we will typically want our arguments to be sound in the double-barrelled sense. But not always! For we sometimes want to argue, and argue absolutely compellingly, from premisses that we *don't* believe to be all true. For example, we might aim to deduce some obviously false consequence from premisses which we don't accept, in the hope that a disputant will agree that this consequence has to be rejected and so come to agree with us that the premisses aren't all true. Sometimes we even want to argue compellingly from premisses we believe are actually inconsistent, precisely in order to bring out their inconsistency. We'll see some examples soon.

(c) As a reality check, note three easy consequences of our definition of soundness:

(1) any sound argument has a true conclusion;
(2) no pair of sound arguments can have conclusions inconsistent with each other;
(3) no sound argument has inconsistent premisses.

Why do these claims hold? For the following reasons:

(1′) A sound argument starts from actually true premisses and involves a necessarily truth-preserving inference move – so it must end up with an actually true conclusion.
(2′) Since a pair of sound arguments will have a pair of actually true conclusions, that means that the conclusions are true together. If they actually *are* true together, then of course they *can* be true together. And if they can be true together then (by definition) the conclusions are consistent with each other.
(3′) Since inconsistent premisses cannot all be true together, an argument starting from those premisses cannot satisfy the first of the conditions for being sound.

Note though that if we replace 'sound' by 'valid' in (1) to (3), we get falsehoods instead. More about valid arguments with inconsistent premisses in due course.

(d) What about arguments where there are intermediate inference steps between the initial premisses and the final conclusion (after all, real-life arguments very often have more than one inference step)? When should we say that they are deductively cogent?

An obvious first shot is to say that such extended arguments are valid when each inference step along the way is valid. And that is part of the story. But we will see in §4.4 that more needs to be said. So we'll hang fire on the question of deductive cogency for multi-step arguments until then.

2.5 'Valid' vs 'true'

Let's pause for a brief terminological sermon, important enough to be highlighted!

> The propositions that occur in arguments as premisses and conclusions are assessed for *truth/falsity*. Inference steps in arguments are assessed for *validity/invalidity*. These dimensions of assessment, as we have stressed, are fundamentally different. We should therefore keep the distinction carefully marked.
>
> So, despite the common misuse of the terms, resolve to *never* again say that a premiss or conclusion or other proposition is 'valid' when you mean it is true. And *never* again say that an argument is 'true' when you mean that it is valid (or sound).

2.6 What's the use of deduction?

(a) Having sharpened our conception of what makes for a valid inference step, let's now say a little more about the central role of such inferences in rational enquiry.

As we noted before, deductively valid inferences are not the only acceptable inferences. Concluding that Jo is injured from the premiss she fell twenty storeys onto concrete is of course perfectly reasonable. Reasoning like this is very often strong enough to trust your life to: the premisses may render the conclusion a racing certainty. But such reasoning isn't deductively valid.

Now consider a more complex kind of inference. Take the situation of the detective, Sherlock let's say. Sherlock assembles a series of clues and then solves the crime by an *inference to the best explanation* (an old term for this is *abductive* reasoning). In other words, the detective arrives at an hypothesis that neatly accommodates all the strange events and bizarre happenings. In the ideally satisfying detective story, this hypothesis strikes us (once revealed) as obviously giving the right explanation – why didn't we think of it? Thus: why is the bed bolted to the floor so it can't be moved? Why is there a useless bell rope hanging by the bed? Why is the top of the rope fixed near a ventilator grille leading through into the next room? What is that strange music heard at the dead of night? All the pieces fall into place when Sherlock infers a dastardly plot to kill the sleeping heiress in her unmovable bed by means of a poisonous snake, trained to descend the rope through the ventilator grille in response to the snake-charmer's music. But although this is an impressive 'deduction' in one everyday sense of the term, it is not deductively valid reasoning in the logician's sense. We may have a number of clues, and the detective's hypothesis H may be the only plausible explanation we can find: but in the typical case it won't be a contradiction to suppose that, despite the way all the evidence stacks up, hypothesis H is actually false. That won't be an inconsistent supposition, only perhaps a very unlikely one. Hence, the detective's plausible 'deductions' are not (normally) valid deductions in the logician's sense.

Now, if our inductive reasoning about the future on the basis of the past is not deductive, and if inference to the best explanation is not deductive either, you might well ask: just how interesting is the idea of deductively valid reasoning? To make the question even more worrisome, consider that paradigm of systematic rationality, scientific reasoning. We gather data, and try to find the best theory that fits; rather like the detective, we aim for the best explanation of the actually observed data. But a useful theory goes well beyond merely summarizing the data. In fact, it is precisely because the theory goes beyond what is strictly given in the data that the theory is useful for making novel predictions. Since the excess content isn't guaranteed by the data, however, the theory cannot be validly deduced from observation statements. So again you might ask: if deductive inference doesn't feature even in the construction of scientific theories, why is it particularly important or interesting? (Yes, it might be crucial for mathematicians – but what about the rest of us?)

(b) But that's far too quick! It is true that we can't simply deduce a scientific theory from the data it is based on. However, it doesn't at all follow that deductive reasoning plays no essential part in scientific reasoning.

Here's a picture of what goes on in science. Inspired by patterns in the data, or by models of the underlying processes, or by analogies with other phenomena, etc., we conjecture that a certain theory is true. Then we use the theory (together with assumptions about 'initial conditions', etc.) to deduce a range of testable predictions.

The first stage, the conjectural stage, may involve flair and imagination, rather than brute logic, as we form our hypotheses. But at the second stage, having made our conjectures, we need to infer testable consequences; and now *this* does involve deductive logic. For we need to examine what else must be true if the hypothesized theory is assumed true: we want to know what our theory deductively entails. Then, once we have deduced testable predictions, we can seek to check them out. Often our predictions prove to be false. We have to reject the theory – or else we have to revise it to accommodate the new data, and then go on to deduce more testable consequences. The process is typically one of repeatedly improving and revising our hypotheses, *deducing* consequences which we can test, and then refining the hypotheses again in the light of test results.

This so-called *hypothetico-deductive* model of science (which highlights the role of deductive reasoning *from theory to predictions*) no doubt needs a lot of development and amplification and refinement. But with science thus conceived, we can see why deduction is absolutely central to the enterprise after all.

And what goes for science, narrowly understood, goes for rational enquiry more widely: deductive reasoning may not be the whole story, but it is an ineliminable core. That's why logic, which teaches us how to appraise passages of reasoning for deductive validity, matters.

2.7 An illuminating circle?

Aristotle wrote in his *Prior Analytics* that "a deduction is speech (*logos*) in which, certain things having been supposed, something . . . results of necessity because of their being so". Our own definition – an argument is valid just if there is no possible situation in which its premisses would be true and its conclusion false – obviously picks up on this key idea that correct deductions are necessarily truth-preserving. This pretty standard definition characterizes what is often called the *classical* concept of validity. And we have tried to say enough to convey an initial understanding of this core idea, and to give a sense of the central importance of deductively valid reasoning.

Note, though, that – as we explained how to understand our definition – we in effect went round in a rather tight circle of interconnected ideas. We first defined deductive validity in terms of what is possible – where we are to understand 'possible' in the weakest, most inclusive, sense. We then further elucidated this relevant weak notion of possibility in terms of what is coherently conceivable. But what is it for a situation to be coherently conceivable? A minimal condition would seem to be that a story about the conceived situation must not involve some hidden self-contradiction. Which presumably is to be understood, in the end, as meaning that we can not validly deduce a contradiction from the story. So now it seems that we have defined deductive validity in a way that needs to be explained, at least in part, by invoking the notion of a valid deduction. How concerning is this?

Well, this kind of circularity is often unavoidable when we are trying to elucidate some really fundamental web of ideas. Often, the best we can do is start with a rough-and-ready, partial, grasp of the various ideas, and then aim to draw out their relationships and make distinctions, clarifying and sharpening the ideas as we explore the web of interconnected notions – so going round, we hope, in an *illuminating* circle. In the present case, this is what we have aimed to do as we have tried to illustrate and explain the notion of deductive validity and interrelated notions. We have at least said enough, let's hope, to get us started and to guide our investigations over the next few chapters.

Still, notions of necessity and possibility do remain genuinely puzzling. So – before you worry that we are starting to build the house of logic on shaky foundations – we should highlight that, looking ahead,

> We will later be giving sharp technical definitions of notions of validity for various special classes of argument, definitions which do *not* directly invoke the notions of necessity/possibility.

These definitions will, however, remain recognizably in the spirit of our preliminary elucidations of the classical concept.

2.8 Summary

> Inductive arguments from past to future, and inferences to the best explanation, are not deductive; but the hypothetico-deductive picture shows why there can still be a crucial role for deductive inference in scientific and other empirical enquiry.
>
> Our preferred definition of the conception of deductive validity: an inference step is deductively valid if and only if there is no possible situation in which its premisses are true and the conclusion false. Call this the classical conception of validity.
>
> The notion of possibility involved in this definition is the weakest and most inclusive.
>
> Equivalently, a deductively valid inference step is necessarily truth preserving, in the strongest sense of necessity.
>
> A one-step argument is valid if and only if its inference step is valid. An argument which is valid and has true premisses is said to be sound.

Exercises 2

(a) Which of the following claims are true and which are false? Explain why the true claims hold good, and give counterexamples to the false claims.

(1) The premisses and conclusion of an invalid argument must together be inconsistent.

(2) If an argument has false premisses and a true conclusion, then the truth of the conclusion can't really be owed to the premisses: so the argument cannot really be valid.

(3) Any inference with actually true premisses and a true conclusion must be truth-preserving and so valid.

(4) You can make a valid argument invalid by adding extra premisses.

(5) You can make a sound argument unsound by adding extra premisses.

(6) You can make an invalid argument valid by adding extra premisses.
(7) If some propositions are consistent with each other, then adding a further true proposition to them can't make them inconsistent.
(8) If some propositions are jointly inconsistent, then whatever propositions we add to them, the resulting propositions will still be jointly inconsistent.
(9) If some propositions are jointly consistent, then their denials are jointly inconsistent.
(10) If some propositions are jointly inconsistent, then we can pick any one of them, and validly infer that it is false from the remaining propositions as premisses.

(b) Show that

(1) If A entails C, and C is equivalent to C', then A entails C'.
(2) If A entails C, and A is equivalent to A', then A' entails C.
(3) If A and B entail C, and A is equivalent to A', then A' and B entail C.

Can we therefore say that 'equivalent propositions behave equivalently in arguments'?

3 Forms of inference

We saw in the first chapter how arguments come in families which share the same type or pattern or form of inference step. Evaluating this shareable form of inference for reliability will then simultaneously give a verdict on a whole range of arguments depending on the same sort of inferential move. In this chapter, we say a little more about the idea of forms of inference and the schemas we use to display them.

3.1 More forms of inference

(a) Consider again the argument:

A (1) No Welshman is a great poet.
(2) Shakespeare is a Welshman.
So (3) Shakespeare is not a great poet.

This is deductively valid, and likewise for the parallel 'Clinton' and 'Bush' arguments which we stated in §2.3. The following is valid too:

A′ (1) No three-year old understands quantum mechanics.
(2) Daisy is three years old.
So (3) Daisy does not understand quantum mechanics.

We can improvise endless variations on the same theme. And plainly, the inference steps in these arguments aren't validated by anything especially to do with poets, presidents or three-year-olds. Rather, they are all valid for the same reason, namely the meaning of 'no' and 'not' and the way that these logical notions distribute in the premisses and conclusion (the same way in each argument). In fact,

Any inference step of the form

No F is G
n is F
So: n is not G

is deductively valid.

As before (§1.5), '*n*' stands in for a name in the italicized schema, while '*F*' and '*G*' (with or without an 'is') stand in for expressions that sort things into kinds or are used to attribute properties. So, substitute as appropriate for the schematic letters (as we can naturally call them), smooth out the language as necessary, and we'll get a valid argument.

That's rather rough. If we want to be more careful, we can put the underlying principle or rule of inference here like this:

> From a pair of propositions, one saying that nothing of some given kind has a certain property, the other saying that a certain individual is of the given kind, we can validly infer the conclusion that the individual in question lacks the property in question.

But it is surely *much* more immediately understandable to use our rough symbolic shorthand (when describing patterns in arguments framed in English). And let's be clear, there is nothing essentially mathematical involved in our use of symbols like '*F*' and '*n*' here. We are simply exploiting the fact that it is easier to talk about the skeletal shape of an inference by *displaying* the relevant pattern using a schema with symbols, instead of trying to *describe* the pattern in cumbersome words.

(b) Here is another argument which we have also met before (§1.4):

B (1) Some philosophy students admire all logicians.
(2) No philosophy student admires any rotten lecturer.
So (3) No logician is a rotten lecturer.

Do the premisses here deductively entail the conclusion?

Consider any situation where the premisses are true. Then by B(1) there will be some philosophy students who admire all logicians. Pick one, Jill for example. We know from B(2) that Jill (since she is a philosophy student) doesn't admire rotten lecturers. That is to say, people admired by Jill aren't rotten lecturers. So in particular, logicians – who are all admired by Jill – aren't rotten lecturers. Which establishes B(3), and shows that the inference step is deductively valid.

What about this next argument?

B′ (1) Some opera fans buy tickets for every new production of Wagner's *Siegfried*.
(2) No opera fan buys tickets for any merely frivolous entertainment.
So (3) No new production of Wagner's *Siegfried* is a merely frivolous entertainment.

This too is valid; and a moment's reflection shows that it essentially involves the same pattern of valid inference step as before.

We can again display the general principle in play using a schema, as follows:

> Any inference step of the following type
>
> *Some F are R to every G*
> *No F is R to any H*
> *So: No G is H*
>
> is deductively valid.

Here we are using '*is/are R to*' to stand in for a form of words expressing a *relation* between things or people. So it might stand in for e.g. 'is married to', 'is taller than', 'is to the left of', 'is a member of', 'loves', 'admires', 'buys a ticket for'. Abstracting from the details of the relations in play to leave just the schematic pattern, we can see that the arguments **B** and **B′** are deductively valid for the same reason, by virtue of sharing the same indicated pattern of inference.

And just try describing the shared pattern of inference *without* using the symbols: it can be done, to be sure, but at what a cost in ease of understanding!

(c) We have now used schemas to display three different types of valid inference steps. We initially met the form of inference we can represent by the schema

All F are G
n is F
So: n is G.

And we have just noted the forms of inference

No F is G
n is F
So: n is not G

Some F are R to every G
No F is R to any H
So: No G is H.

Any argument instantiating one of these patterns will be deductively valid. The way that 'all', 'every' and 'any', 'some', 'no' and 'not' distribute between the premisses and conclusion means that inferences following these schematic forms are necessarily truth-preserving – if the premisses are true, the conclusion has to be true too.

For the moment, let's just note three more examples of deductively reliable forms of inference, involving different numbers of premisses:

No F is G
So: No G is F

All F are H
No G is H
So: No F is G

All F are either G or H
All G are K
All H are K
So: All F are K.

(Convince yourself that arguments which instantiate these schemas are indeed valid in virtue of the meanings of the logical words 'all', 'no' and 'either . . . or . . . '.)

3.2 Four simple points about the use of schemas

The trick of using schemas to display patterns of reasoning is not only used in the logic classroom: it is, for example, quite common in philosophical writing to use schemas in a rough-and-ready way in order to clarify the structure of arguments.

Later, from Chapter **??** on, we use similar schemas in a significantly more disciplined way in talking about patterns of reasoning in arguments framed in artificial formal languages – languages that logicians love, for reasons which will soon become clear.

We pause next, then, to emphasize four initial points that apply to both the informal and the more formal uses of schemas to display types of inference step. (These should be very obvious points: but it is worth reading them into the record, to block any possible misunderstandings.)

(a) Take again the now familiar form of inference we can display like this:

All F are G
n is F
So: n is G.

Does the following argument count as an instance of this schematic pattern?

C (1) All men are mortal.
(2) Tweety is a robin.
So (3) Donald Trump is female.

Well, of course not! But let's spell out why.

True enough, C(1) attributes a certain property to everything of a given kind; so taken by itself, we can informally represent it as having the shape *All F are G*. Likewise, C(2) and (C3) are simple propositions which, taken seperately, can be represented as having the form *n is F* or *n is G*. However, when we describe an inference as an instance of the displayed schema we are indicating that the property-ascribing term *F* is involved in the two premisses. Likewise, we are indicating that the name *n* and general term *G* that occur in the premisses recur in the conclusion. This is worth highlighting:

> The whole point of using recurring symbols in schemas representing forms of inference is to indicate patterns of recurrence in the premisses and conclusion. Therefore, when we fill in a schema by substituting expressions for the symbols, we must follow the rule: same symbols, same substitutions.

So, in the present case, in moving back from our displayed abstract schema to a particular instance of it, we must preserve the patterns of recurrence by being consistent in how we fill substitute for the '*F*'s and '*G*'s and '*n*'s.

(b) What about the following argument? Does this also count as an instance of the last schema we displayed?

D (1) All men are men.
(2) Socrates is a man.
So (3) Socrates is a man.

Instead of filling in our schema at random, we have at least this time been consistent, substituting uniformly for both occurrences of '*F*', and likewise for both occurrences of '*G*'. However, we happen to have substituted in the same way each time.

We will allow this. So, to amplify, the rule is: same schematic letter, same substitutions – but not different schematic letters, necessarily different substitutions. In the present case, if the '*F*'s and '*G*'s are both substituted in the same way, we still get a valid argument – **D** obviously cannot have true premisses and a false conclusion!

Of course, **D** is no use at all as a means for persuading someone of the conclusion – to use the argument to establish its conclusion you'd already need to accept that very same proposition as a premiss. However, there is a very important distinction that needs reinforcing here. Being a valid, or even a sound, argument with a necessarily truth-preserving inference step is one thing, being usefully persuasive is something else (an argument won't persuade you of the truth of the conclusion if you don't actually believe the premisses, and it won't be useful if you already have to believe the conclusion in order to believe the premisses).

Keeping in mind that distinction between being valid and being usefully persuasive, counting **D** as a limiting case of a deductively virtuous argument is acceptable. After all, an inference step that covers no ground has no chance to go wrong.

(c) And what about the following argument? How does this stand with respect to our displayed schema?

E (1) Socrates is a man.
(2) All men are mortal.
So (3) Socrates is mortal.

Compared with the abstract schema, the premisses here are stated 'in the wrong order', with the general premiss second. But so what? An inference step is valid, we said, just when there is no possible situation in which its premisses are true and its conclusion false. Given that definition, the validity of the inference doesn't depend at all on the order in which the premisses happen to be stated. Hence for our purposes – when we want to display the form of an inference – the order in which the various premisses are represented in a schema is irrelevant. Hence we can take **E** as exemplifying just the same pattern of inference as before.

(d) Last but not least, what is the relation between our symbolic schema at the beginning of this section and its three variants below?

All F are G	*All H are K*	*All* Φ *are* Ψ	*All* ① *are* ②
n is F	*m is H*	*α is* Φ	⋆ *is* ①
So: n is G	*So: m is K*	*So: α is* Ψ	*So:* ⋆ *is* ②

Plainly, these four *different* schemas are alternative possible ways of representing the *same* form or pattern of inference (so you must not confuse schemas with forms of inference!) Remember: what we are trying to reveal is a pattern of recurrence. The recurrent '*m*'s and '*α*'s and '⋆'s are just different arbitrary ways of representing how names will recurring in the common pattern. Likewise the '*F*'s or 'Φ's or whatever are just different arbitrary ways of representing how repeated property-ascribing terms are to be substituted. We can use whatever symbols take our fancy to do the job. Later we will in fact mainly be using Greek letters; but for our current informal purposes we will stick to ordinary italic letters.

3.3 Arguments can instantiate many patterns

(a) It can be tempting to talk about 'the' form or pattern of inference exemplified by an argument. But we do need to be *very* careful here.

Consider again the argument **E** from the two premisses *All men are mortal*, and *Socrates is a man*, to the conclusion *Socrates is mortal*. It is, trivially, an instance of the wildly unreliable form of inference

(1) *A, B, so C*

where '*A*' etc. stand in for whole propositions, complete premisses and conclusions.

Exposing some structure in the individual propositions, but still ignoring the all-important recurrences, **E** also instantiates the equally unreliable form of inference

(2) *All F are G, m is H, so n is K*

since it has the right sort of general premiss, a simple second premiss ascribing some property to a named individual, and a similar conclusion.

Next, and now filling in enough details about the structure of the inference to bring out an inferentially reliable pattern of recurrence, the 'Socrates' argument also instantiates, as we said before, the form

(3) *All F are G, n is F, so n is G.*

But we can go further. **E** is also an instance of e.g. the perfectly reliable types of inference

(4) *All F are G, Socrates is F, so Socrates is G,*

(5) *All men are G, n is a man, so n is G.*

And we could even, going to the extreme, take the inference to be the one and only example of the reliable type

(6) *All men are mortal, Socrates is a man, so Socrates is mortal,*

which is (so to speak) a pattern with all the details filled in!

In sum, the 'Socrates' argument **E** exemplifies a number of different forms or patterns of inference, specified at different levels of generality. Which all goes to show:

> There is no such thing as *the* unique form or pattern of inference that a given argument can be seen as instantiating.

Of course, this isn't at all to deny that the form of inference (3) has a special place in the story about **E**. For (3) is the most general but still reliable pattern of inference that the argument exemplifies. In other words, the schema *All F are G; n is F; so, n is G* gives just enough of the structure of **E** – but no more than we need – to enable us to see that the argument is in fact valid. Further, it is the reliability of *this* pattern of inference which someone defending the argument **E** will surely appeal to. It is rather natural, then, to fall into talking of the schema as revealing 'the' pattern of inference in **E** – but, as we've just seen, that would be loose talk.

(b) A terminological note:

> It is common to refer to a universally reliable form or pattern of inference – i.e. one whose instances are all valid – as itself being 'valid'.

But if we fall in with this usage, it is very important to emphasize again that, as we have just illustrated, an 'invalid' pattern of inference like (1) or (2) – meaning a pattern whose instances are not *all* valid – can still have *some* instances that do happen to be valid (being valid, of course, for some other reason than exemplifying the unreliable pattern).

3.4 Summary

Forms or patterns of inference can be conveniently represented by schemas using symbolic letters, whose use is governed by the obvious convention that a symbol stands in for the same expression whenever it appears within an argument schema.

There is strictly no such thing as the unique pattern or form exemplified by an argument. Valid arguments will typically be instances of various reliable forms; but they can also be instances of other, 'too general', unreliable forms.

A form of inference is said to be valid if all its instances are valid.

Exercises 3

(a) Which of the following patterns of inference are deductively reliable, meaning that all their instances are valid? If you suspect an inference pattern is unreliable, find an instance which obviously fails because it has plainly true premisses and a false conclusion.

(1) Some *F* are *G*; no *G* is *H*; so, some *F* are not *H*.
(2) Some *F* are *G*; some *F* are *H*; so, some *G* are *H*.
(3) All *F* are *G*; some *F* are *H*; so, some *H* are *G*.
(4) No *F* is *G*; some *G* are *H*; so, some *H* are not *F*.
(5) No *F* is *G*; no *G* is *H*; so, some *F* are not *H*.
(6) All *F* are *G*; no *G* is *H*; so, no *H* is *F*.

Arguments of these kinds, with two premisses and a conclusion, with each proposition being of one of the four kinds 'All . . . are . . . ', 'No . . . is . . . ', 'Some . . . are . . . ' or 'Some . . . are not . . . ', with three general terms filling in the gaps (the two terms in the conclusion occurring in separate premisses), are the traditional categorical *syllogisms* first discussed by Aristotle. Which valid types of syllogism of this kind have a conclusion of the form 'All *F* are *H*'? Which have a conclusion of the form 'Some *F* are not *H*'?

(b) What of the following patterns of argument? Are these deductively reliable?

(1) All *F* are *G*; so, nothing that is not *G* is *F*.
(2) All *F* are *G*; no *G* are *H*; some *J* are *H*; so, some are not *F*.
(3) There is an odd number of *F*, there is an odd number of *G*; so there is an even number of things which are either *F* or *G*.
(4) All *F* are *G*; so, at least one thing is *F* and *G*.
(5) *m* is *F*; *n* is *F*; so, there are at least two *F*.
(6) All *F* are *G*; no *G* are *H*; so, all *H* are *H*.

4 Proofs

In Chapter 1, we explained what it is for an inference to be deductively valid, and then we noted that different arguments can share a common form or pattern of inference. Chapters 2 and 3 explored these ideas a little further.

We now move on to consider the following question: how can we establish that a not-obviously-valid inference is in fact valid? One answer, in headline terms, is: by giving a multi-step argument that takes us from that inference's premisses to its conclusion by simple, obviously valid, steps – in other words, by giving a derivation or *proof*. This chapter explains.

4.1 Proofs: first examples

(a) Take a charmingly silly example from Lewis Carroll:

A Babies are illogical; nobody is despised who can manage a crocodile; illogical persons are despised; so babies cannot manage a crocodile.

This three-premiss, one-step, argument is in fact valid. That is probably already obvious. But how can we demonstrate the argument's validity if it isn't immediately clear? Well, consider the following two-step derivation:

A′	(1)	Babies are illogical.	(premiss)
	(2)	Nobody is despised who can manage a crocodile.	(premiss)
	(3)	Illogical persons are despised.	(premiss)
	(4)	Babies are despised.	(from 1, 3)
	(5)	Babies cannot manage a crocodile.	(from 2, 4)

Here, we have inserted an extra step between the original premisses and the target conclusion (and we can omit writing 'So' before lines (4) and (5), as the commentary on the right suffices to indicate that they are the results of inferences).

The inference from the original premisses to the interim conclusion (4), and then the further inference from (one of) the original premisses plus that interim claim (4) to the final conclusion (5), should both be evidently valid. We can therefore get from the initial premisses to the final conclusion by a necessarily truth-preserving route. Which shows that the inferential leap in the original argument **A** is indeed valid.

(b) Here's a second quick example, equally daft (borrowed from Richard Jeffrey):

B Everyone loves a lover; Romeo loves Juliet; so everyone loves Juliet.

Take the first premiss to mean 'everyone loves anyone who is a lover' (where a lover is, by definition, a person who loves someone). Then this inference too is deductively valid! Here is a multi-step derivation:

B′	(1)	Everyone loves a lover.	(premiss)
	(2)	Romeo loves Juliet.	(premiss)
	(3)	Romeo is a lover.	(from 2)
	(4)	Everyone loves Romeo.	(from 1, 3)
	(5)	Juliet loves Romeo.	(from 4)
	(6)	Juliet is a lover.	(from 5)
	(7)	Everyone loves Juliet.	(from 1, 6)

We have again indicated on the right the 'provenance' of each new statement as the argument unfolds. And by inspection we can see that each small inference step is valid, i.e. is necessarily truth-preserving. As the argument grows, then, we are adding new propositions which must also be true, assuming the original premisses are true. These new true propositions can then serve in turn as inputs to further valid inferences, yielding more truths. Hence everything is chained together so that, if the original premisses are true, each added proposition must be true too, and therefore the final conclusion in particular must be true. Hence the original inference in **B** that jumps straight across the intermediate stages must be valid.

(c) These very simple examples illustrate, then, the standard technique for establishing the validity of an unobviously valid inference step – the idea goes back to Aristotle.

> One way of demonstrating that an inferential leap from some premisses to a given conclusion is deductively valid is by breaking down the big leap into smaller inference steps, each one of which is clearly deductively valid.

There is a familiar term for a chain of inferences put together in such a way as to be deductively cogent – it's a *proof*. So the idea is that we can establish that an inference step is valid by providing a proof filling in between the original premisses and conclusion, where each step in the proof is evidently in good order. (Careful though: a proof in this sense doesn't necessarily establish the truth of its conclusion outright; it only shows the truth of the conclusion assuming the truth of the premisses.)

4.2 Fully annotated proofs

(a) In our first two examples of proofs, we have indicated at each new step which earlier statements the inference depended on. But we can do even better by also explicitly indicating what type or *form* of inference move is being invoked at each stage.

For use in our next example, then, recall two principles that we have met before:

Any inference step of either of the following two types is valid,

(R) *No F is G*
n is F
So: n is not G

(S) *No F is G*
So: No G is F

And now consider the following little argument:

C No logician is bad at abstract thought. Jo is bad at abstract thought. So Jo is not a logician.

This hardly needs a proof! – but still, we *can* prove it, invoking (R) and (S), as follows:

C′	(1)	No logician is bad at abstract thought.	(premiss)
	(2)	Jo is bad at abstract thought.	(premiss)
	(3)	No-one bad at abstract thought is a logician.	(from 1 by S)
	(4)	Jo is not a logician.	(from 3, 2 by R)

(Why do you think we have mentioned steps (3) and (2) in that order in the annotation for line (4)?)

Note that in treating the move from (1) to (3) as an instance of (S) we are – strictly speaking – cutting ourselves a bit of slack, quietly massaging the English grammar and inserting an 'a'. (It's actually rather difficult to give strict rules for doing this sort of thing, even it causes us no difficulty in practice: which is one good reason for eventually moving from considering arguments framed in messy English to considering their counterparts in a nicely rule-governed formalized language.)

(b) Of course, the validity of **C** has nothing especially to do with Jo, logicians and abstract thought. We can say, more generally:

Any inference step of the following type is valid,

(T) *No F is G*
n is G
So: n is not F

Further, any instance of this pattern can be shown to be valid by using a proof parallel to **C′**. In other words, the reliability of (T) follows from the reliability of (R) and (S).

Which is the point of mentioning this little example. It reveals another important sense in which logical enquiry can be systematic (compare §1.5). Not only can we treat arguments wholesale by considering together whole families relying on the same principles of inference, but we can systematically interrelate patterns of inference like (R), (S) and (T) by showing how the reliability of some ensures the reliability of others.

So logical theory will give us much more than an unstructured catalogue of deductively reliable principles of inference; we can and will explore how principles of inference hang together. More about this later in the book!

(c) Let's take another argument from Lewis Carroll, who is an inexhaustible source of very silly examples (but there is no point in getting too solemn at this stage!):

D Anyone who understands human nature is clever; every true poet can stir the heart; Shakespeare wrote Hamlet; no one who does not understand human nature can stir the heart; no one who is not a true poet wrote Hamlet; so Shakespeare is clever.

The big inferential leap from those five premisses to the conclusion is in fact valid. For note the following:

Any inference step of either of the following two types is valid:

(U) *Any/every F is G*
n is F
So: n is G

(V) *No one who isn't F is G*
So: Any G is F

Then we can argue as follows, using these two forms of valid inference step:

D′

(1)	Anyone who understands human nature is clever.	(premiss)
(2)	Every true poet can stir the heart.	(premiss)
(3)	Shakespeare wrote Hamlet.	(premiss)
(4)	No one who does not understand human nature can stir the heart.	(premiss)
(5)	No one who is not a true poet wrote Hamlet.	(premiss)
(6)	Anyone who wrote Hamlet is a true poet.	(from 5, by V)
(7)	Shakespeare is a true poet.	(from 3, 6, by U)
(8)	Shakespeare can stir the heart.	(from 2, 7, by U)
(9)	Anyone who can stir the heart understands human nature.	(from 4, by V)
(10)	Shakespeare understands human nature.	(from 8, 9, by U)
(11)	Shakespeare is clever.	(from 1, 10, by U)

That's all very *very* laborious: but now we have the provenance of every move fully annotated. Each of the inference steps is an instance of a clearly truth-preserving pattern, and together they get us from the initial premisses of the argument to its final conclusion. Hence, if the initial premisses are true, the conclusion must be true too. So our proof here really does show that the original inference **D** is deductively valid.

4.3 Glimpsing an ideal

(a) Using toy examples, we have now glimpsed an ideally explicit way of setting out one kind of regimented multi-step proof. Every premiss we are going to need gets stated at the outset and marked as such; and then we write further propositions one under another, each coming with a certificate which tells us what it is inferred from and what principle of inference is being invoked.

Needless to say, everyday arguments (and even not-so-everyday proofs in mathematics books) normally fall a very long way short of meeting these ideal standards of explicitness! They rarely come ready-chunked into numbered statements, with each new inference step bearing a supposed certificate of excellence. However, faced with a puzzling multi-step argument, massaging it into something nearer this fully documented shape will help us to assess the argument. On the one hand, an explicitly annotated proof gives any critic a particularly clear target to fire at. If someone wants to reject the conclusion, then they will have to rebut one of the premisses, show an inference step doesn't have the claimed form, or come up with a counterexample to the reliability of one of the forms of inference that is explicitly called on. On the other hand, looking on the bright side, if the premisses are agreed and if the inference moves are also agreed to be uncontentiously valid, then the proof will establish the conclusion beyond further dispute.

(b) The exercise of regimenting an argument into a more ideally explicit form will also reveal *redundancies* (premisses or inferential detours that are not really needed to get to the final conclusion) and expose where there are *suppressed premisses* (i.e. premisses that are not actually stated but which are needed if we are to get a well-constructed proof).

Traditionally, arguments with suppressed premisses are called *enthymemes*. Without worrying about that label, here's a toy example:

E The constants of nature have to take values in an extremely narrow range (have to be 'fine-tuned') to permit the evolution of intelligent life. So the universe was intelligently designed.

As it stands this is plainly gappy: there are unspoken assumptions in the background. But it is instructive to try to turn it into a deductively valid argument, by making explicit the currently suppressed premisses. One such premiss will be uncontroversial, just noting that intelligent life has evolved. The other premiss will be much more problematic: roughly, intelligent design is needed in order for the universe to be fine-tuned enough. Only with some such additions can we get a deductively cogent argument.

4.4 Deductively cogent multi-step arguments

When is a multi-step argument internally cogent? The answer is predictable: when each inference step is cogent and also the steps are chained together in the right kind of way. But it is worth spelling out these two requirements a bit more carefully.

(a) So first consider the following mini-argument:

F	(1) All philosophers are logicians.	(premiss)
	(2) All logicians are philosophers.	(from 1?!)

This inference is horribly fallacious. It just doesn't follow from the premiss that all philosophers are logicians that only philosophers are logicians. (You might as well argue 'All women are human beings, hence all human beings are women'.)

Here's another really bad argument:

G	(1) All existentialists are philosophers.	(premiss)
	(2) All logicians are philosophers.	(premiss)
	(3) All existentialists are logicians.	(from 1, 2?!)

It plainly doesn't follow from the claims that the existentialists and logicians are both among the philosophers that any of the existentialists are logicians, let alone that all of them are. (You might as well argue 'All women are human beings, all men are human beings, hence all women are men'.)

But now imagine that someone chains this pair of rotten inferences together into a two-step argument as follows:

H (1) All existentialists are philosophers. (premiss)
(2) All philosophers are logicians. (premiss)
(3) All logicians are philosophers. (from 2?!)
(4) All existentialists are logicians. (from 1, 3?!)

Our reasoner first makes the same fallacious inference as in **F**; and then they compound the sin by committing the same howler as in **G**. So they have gone from the initial premisses H(1) and H(2) to their final conclusion H(4) by two quite terrible moves.

Yet note that in this case – despite the howlers along the way – the final conclusion H(4) in fact really does follow from the initial premisses H(1) and H(2). If the existentialists are all philosophers, and all philosophers are logicians, then the existentialists must of course be logicians!

What's the moral of this example? If we were to say that a multi-step argument is deductively cogent just if the big leap from the initial premisses to the final conclusion is valid, then we'd have to count the two-step argument **H** as a cogent proof. Which would be a *very* unhappy way of describing the situation, given that the two-step derivation involves a couple of nasty fallacies!

For a genuine proof – a deductively cogent multi-step argument, where everything is in good order – we should therefore require not just that the big leap from the initial premisses to the final conclusion is valid, but that the individual inference steps along the way are all valid too. Exactly as you would expect!

(b) But that is not the whole story. For deductive cogency, we also need a proof's valid inference steps to be chained together in the right kind of way.

To illustrate, consider the inference step here:

I (1) Socrates is a philosopher. (premiss)
(2) All philosophers have snub noses. (premiss)
(3) Socrates is a philosopher and all philosophers have snub noses. (from 1, 2)

That's trivially valid (since from *A* together with *B* you can infer *A-and-B*).

Next, here is another equally trivial valid inference (since from *A-and-B* you can of course infer *B*):

J (1) Socrates is a philosopher and all philosophers have snub noses. (premiss)
(2) All philosophers have snub noses. (from 1)

And thirdly, here is another plainly valid inference of a now very familiar form:

K (1) Socrates is a philosopher. (premiss)
(2) All philosophers have snub noses. (premiss)
(3) Socrates has a snub nose (from 1, 2)

(Note that **I** and **K** have the same premisses and different conclusions – but that's fine. We can extract more than one implication from the same premisses!)

Taken separately, then, those three little inferences are entirely unproblematic. However, imagine now that someone chains these inferences together to get the following unholy tangle:

L	(1)	Socrates is a philosopher.	(premiss)
	(2)	Socrates is a philosopher and all philosophers have snub noses.	(from 1, 3, as in **I**)
	(3)	All philosophers have snub noses.	(from 2 as in **J**)
	(4)	Socrates has a snub nose	(from 1, 3 as in **K**)

By separately valid steps we seem to have deduced the shape of Socrates' nose just from the premiss that he is a philosopher! What has gone wrong?

The answer is plain. In the middle of the argument we have gone round in a circle. L(2) is derived from inputs including L(3), and then L(3) is derived from L(2). Circular arguments like this can't take us anywhere. Which shows that there is another condition for being a deductively cogent multi-step argument. As well as the individual inference steps all being valid, the steps must be chained together in a non-circular way. In particular, each new step must depend only on what's gone before.

4.5 Indirect arguments

(a) So far, so straightforward. But now we need to complicate the story, introducing an important new idea. The headline news: we often aim to establish a conclusion *not-S* by supposing the opposite *S*, and showing that this leads to something absurd.

Return to the silly crocodile argument **A′**. Imagine that someone has a momentary lapse and pauses over the final step (where we inferred *Babies cannot manage a crocodile* from the original premiss *Nobody is despised who can manage a crocodile* together with the interim conclusion *Babies are despised*). How could we convince them that this step really is valid?

We might try amplifying the argument so we get an expanded (not quite fully documented) proof which now runs as follows:

A″	(1)	Babies are illogical.	(premiss)
	(2)	Nobody is despised who can manage a crocodile.	(premiss)
	(3)	Illogical persons are despised.	(premiss)
	(4)	Babies are despised.	(from 1, 3)
		Suppose temporarily, for the sake of argument,	
	(5)	Babies *can* manage a crocodile.	(supposition)
	(6)	Babies are not despised.	(from 2, 5)
	(7)	Contradiction!	(from 4, 6)
		Our supposition leads to absurdity, hence	
	(8)	Babies cannot manage a crocodile.	(from 5–7 by RAA)

What is going on in this expanded proof? We want to establish (8). But this time, instead of aiming directly for the conclusion, we branch off by temporarily supposing the exact opposite is true, i.e. we suppose (5). However, this supposition immediately leads to something that flatly contradicts an earlier claim. Hence the temporary supposition (5) has to be rejected as leading to absurdity (given the original premisses). Therefore its denial (8) must be true after all.

The terse justification for the final step '(from 5–7 by RAA)' indicates that we have made what is traditionally called a *reductio ad absurdum* inference – an assumption made for the sake of argument at (5) leads to absurdity at (7), so has to be rejected.

Note by the way that, in the indented part of the proof, we are working with our initial premisses plus a new temporary supposition all in play, assumptions which taken together are in fact inconsistent. By the end of the (short!) indented bit of the proof, we have exposed this inconsistency by showing that those assumptions together entail a contradiction. As we said back in §2.4, it is very important that we *can* in this way argue validly from inconsistent assumptions to expose their inconsistency.

(b) Here is another quick example. Take the argument:

M No girl loves any unreconstructed sexist; Caroline is a girl who loves whoever loves her; Henry loves Caroline; hence Henry is not an unreconstructed sexist.

This is valid. Here is one way to see that it is, by using another reductio argument:

M′	(1)	No girl loves any unreconstructed sexist.	(premiss)
	(2)	Caroline is a girl who loves whoever loves her.	(premiss)
	(3)	Henry loves Caroline.	(premiss)
	(4)	Caroline is a girl who loves Henry.	(from 2, 3)
	(5)	Caroline is a girl.	(from 4)
	(6)	Caroline loves Henry.	(from 4)
		Suppose temporarily, for the sake of argument,	
	(7)	Henry is an unreconstructed sexist.	(supposition)
	(8)	No girl loves Henry.	(from 1, 7)
	(9)	Caroline does not love Henry.	(from 5, 8)
	(10)	Contradiction!	(from 6, 9)
		Our supposition leads to absurdity, hence	
	(11)	Henry is not an unreconstructed sexist.	(from 7–10 by RAA)

(c) Let's now spell out the RAA principle invoked in these two proofs (the idea is surely a familiar one, even if we are now spelling it out more explicitly that you are used to). We can put it like this:

> *Reductio ad absurdum* If $A_1, A_2, \ldots, A_n$ (as background premisses) plus the temporary supposition S entail a contradiction, then from $A_1, A_2, \ldots, A_n$ by themselves we can validly infer *not*-S.

Here, the 'A's and 'S' stand in for whole propositions; and entailing a contradiction is a matter of entailing some proposition C while also entailing its denial *not*-C (or equivalently, entailing the single proposition C *and not*-C).

Why does this principle hold? Suppose that the *A*s plus *S* *do* entail a contradiction *C and not-C*. Now, what is entailed by truths must also be true (since entailment is truth-preserving). So suppose that the *A*s plus *S* were all true together. Then *C and not-C* would be true too. But that's ruled out, as a contradiction can never be true. So the *A*s plus *S* cannot all be true together. Of course, just knowing that at least one of *A*s plus *S* must be false doesn't tell us where to pin the blame! However, we do know this: in any situation in which the *A*s *are* all true, the remaining proposition *S* must be the false one. Hence those propositions $A_1, A_2, \ldots, A_n$ entail *not-S*.

(Question: What happens in the case where $n = 0$ and there are no background premisses?)

(d) Reductio ad absurdum arguments are often called *indirect* arguments. We don't go straight from premisses to conclusion, but take a side-step via some additional supposition which we temporarily add for the sake of argument and then eventually 'discharge'. (We have visually signalled the side-step by indenting the column of argument to the right when bring a new supposition into play, and going back left when the supposition is dropped again.)

As we will see later, there are other familiar kinds of indirect arguments. And to cover such indirect forms of reasoning, we will have to extend what we said before about cogent multi-step arguments. In particular, we must now permit steps which introduce new temporary assumptions 'for the sake of argument' (so long as they are clearly flagged). Then, as well as ordinary valid inference steps which depend on what we have previously established and/or some new assumption(s) currently in play, we will also allow inference steps like (RAA), which allow us to draw consequences from 'subproofs' starting with temporary assumptions. Our discussion here necessarily can only give a hint of what is to come: we will return to the topic when we eventually start discussing formal 'natural deduction' proofs in Chapter **??**.

4.6 Summary

To establish the validity of a perhaps unobvious inferential leap we can use deductively cogent multi-step arguments, i.e. proofs, filling in the gap between premisses and conclusion.

Simple, direct, proofs chain together inference steps that are valid – ideally, *obviously* valid – building up from the initial premisses to the desired conclusion, with each new step depending on what's gone before.

However, some common methods of proof like reductio ad absurdum are 'indirect', and involve making new temporary suppositions for the sake of argument, suppositions that are later discharged.

The availability of indirect modes of inference complicates the story about what makes for a deductively cogent multi-step argument; we will return to this.

Exercises 4

Which of the following arguments are valid? Where an argument is valid, provide a proof. Some of the examples are enthymemes that need repair.

(1) No philosopher is illogical. Jones keeps making argumentative blunders. No logical person keeps making argumentative blunders. All existentialists are philosophers. So, Jones is not an existentialist.

(2) Jane has a first cousin. Jane's father had no siblings. So, if Jane's mother had no sisters, she had a brother.

(3) Every event is causally determined. No action should be punished if the agent isn't responsible for it. Agents are only responsible for actions they can avoid doing. Hence no action should be punished.

(4) Something is an elementary particle only if it has no parts. Nothing which has no parts can disintegrate. An object that cannot be destroyed must continue to exist. So an elementary particle cannot cease to exist.

(5) No experienced person is incompetent. Jenkins is always blundering. No competent person is always blundering. So, Jenkins is inexperienced.

(6) Only logicians are good philosophers. No existentialists are logicians. Some existentialists are French philosophers. So, some French philosophers are not good philosophers.

(7) Either the butler or the cook committed the murder. The victim died from poison if the cook did the murder. The butler did the murder only if the victim was stabbed. The victim didn't die from poison. So, the victim was stabbed.

(8) Promise-breakers are untrustworthy. Beer-drinkers are very communicative. A man who keeps his promises is honest. No one who doesn't drink beer runs a bar. One can always trust a very communicative person. So, no one who keeps a bar is dishonest.

(9) When I do an example without grumbling, it is one that I can understand. No easy logic example ever makes my head ache. This logic example is not arranged in regular order, like the examples I am used to. I can't understand these examples that are not arranged in regular order, like the examples I am used to. I never grumble at an example, unless it gives me a headache. So, this logic example is difficult.

5 The counterexample method

Suppose we are unsure whether some inferential move is valid or not. If we can find a multi-step argument from the given premisses to the claimed conclusion, relying on uncontentiously valid inference steps, then that settles it – the inference *is* valid. But what if we have tried to find such a proof and failed? What does that show? Maybe we have just not spotted a proof, although the inference is again valid. Or maybe the inference is invalid. But if failure to find a proof doesn't settle the matter one way or the other, how *can* we demonstrate that a challenged inference really is invalid?

In this chapter we put together the observation that different arguments can depend on the same pattern of inference (as outlined in Chapter 3) with the trivial invalidity principle (which we met in Chapter 2) to give us one simple method for showing that various invalid inferences are indeed deductively invalid.

5.1 'But you might as well argue …'

(a) In fact, we have already quietly used this method in passing in §4.4. We noted that the inference step in the mini-argument

A (1) All philosophers are logicians.
So (2) All logicians are philosophers.

is invalid. That should have been immediately obvious. But, to press home the point, we remarked that you might as well argue

A′ (1) All women are human beings.
So (2) All human beings are women.

What is the force of this brisk comparison? We can unpack it as follows:

> Obviously, argument **A′** has a true premiss and a false conclusion; hence by the invalidity principle of §2.3 it can't be valid. But the inferential move in argument **A** is no better. There is nothing to separate the cases. The arguments evidently stand or fall together as far as validity is concerned. Hence, **A** is invalid too.

And that's a simple illustration of the basic method we need. Roughly: to show an inference step is invalid, find an argument which relies on the same form of inference but which is clearly invalid by the invalidity principle.

(b) We used the same method on a second example in §4.4. But let's next have a new illustration, this time revisiting an argument we met in §1.4:

B (1) Most Irish are Catholics.
(2) Most Catholics oppose abortion.
So (3) At least some Irish oppose abortion.

We quickly persuaded ourselves that the inference here is hopeless by imagining a situation in which the premisses would be true and conclusion false. But equally, we could have highlighted the parallel between **B** and

B′ (1) Most chess grandmasters are men.
(2) Most men are no good at chess.
So (3) At least some chess grandmasters are no good at chess.

What is the force of the comparison again?

> Obviously argument **B′** is invalid. Chess is still (at least at the upper levels of play) a predominantly male activity, though one that few men are any good at. The premisses of argument **B′** are true. But the conclusion is false. So **B′** is invalid by the invalidity principle. But the inferential move in argument **B** is no better. The arguments evidently stand or fall together as far as validity is concerned. Hence, **B** is invalid too.

(c) There is nothing mysterious or difficult or even novel about the 'But you might as well argue . . . ' gambit for showing an inference to be invalid. This is another technique already noted and used by Aristotle, and we use it all the time in the everyday evaluation of arguments.

For example, some gossip says that Mrs Jones must be an alcoholic, because she has been seen going to the Cheapo Booze Emporium and everyone knows that that is where all the local alcoholics hang out. You reply, 'But you might as well argue that Bernie Sanders is a Republican Senator, because he's been seen going into the Senate, and everyone knows that that's where all the Republican Senators hang out'. Which vividly brings out why the gossip's inference won't do.

5.2 The counterexample method, more carefully

(a) Remember, however, that very important point from §3.3. A form of inference may not in general be reliable, it may have many instances J which are invalid, yet it can still have some special instances I which do happen to be valid. So we *can't* say, crudely, 'This inference I is an instance of the form F; but here's another instance J of the same form F which is plainly invalid. So I is invalid too.'

We therefore need to be spell out carefully our 'counterexample method' for showing invalidity. The idea is this:

> *Showing invalidity by counterexample*
>
> *Stage 1* Given a (one-step) argument whose validity is up for assessment, first locate some form of inference that this argument actually depends on – meaning a form of inference that needs to be reliably truth-preserving if the inference in the argument is indeed to be valid.

Stage 2 Show that this form of inference is *not* a generally reliable one by finding a *counterexample*. In other words, find another argument having the same pattern of inference which is uncontroversially invalid, e.g. because it has actually true premisses and a false conclusion so we can apply the invalidity principle.

Note, by the way, that we do not here *require* a counterexample to be generated by the invalidity principle, i.e. to have actually true premisses and an actually false conclusion. Telling counterexamples don't *have* to be constructed from 'real life' situations. A merely imaginable but uncontroversially coherent counterexample will do just as well. Why? Because that's still enough to show that the inferential pattern in question is unreliable: it doesn't necessarily preserve truth in all possible situations. However, using 'real life' counterexamples does often have the great advantage that you don't get into any disputes about what is coherently imaginable.

(b) Obviously, in using this method for showing that a given argument is invalid, everything turns on picking out some form of inference which the challenged argument requires to be reliable.

Often that is very easy to do. Take our initial example 'All philosophers are logicians. So all logicians are philosophers.' What form of inference move can this argument possibly be relying on, other than *All F are G; so all G are F*? And, as we saw, it is trivial to come up with a counterexample to the general reliability of *this* form of inference. Similarly for our second example 'Most Irish are Catholics. Most Catholics oppose abortion. So at least some Irish oppose abortion.' What form of inference can that argument be relying on, other than *Most F are G; most G are H; so at least some F are H*? It is trivial to come up with a counterexample to the validity of that form of inference too.

Other cases, however, can be more problematic. In practice, an exchange can go like this: someone proposes an argument; you challenge it by replying 'But you might as well argue . . . ' (advancing a seemingly parallel but patently invalid argument); but then the proponent of the challenged argument seeks to show that your supposed counterexample isn't a fair one – the original argument didn't actually depend on the pattern of inference you supposed. With arguments served up in ordinary prose, this might not be easy to settle. But, at the very least, challenge-by-apparent-counterexample will force the proponent of an argument to clarify what is supposed to be going on, to make it plain what principle of inference is being used.

5.3 A 'quantifier shift' fallacy

(a) Let's have an example to illustrate the last point. Consider the following quotation (in fact, the opening words of Aristotle's *Nicomachean Ethics*):

> Every art and every enquiry, and similarly every action and pursuit, is thought to aim at some good; and for this reason the good has rightly been declared to be that at which all things aim.

At first sight, there is an argument here with an initial premiss we can sum up as

C (1) Every practice aims at some good.

And then a conclusion is drawn (note the inference marker 'for this reason'):

So (2) There is some good ('*the* good') at which all practices aim.

But this argument has a very embarrassing similarity to the following one:

C′ (1) Every assassin's bullet is aimed at some victim.
So (2) There is some victim at whom every assassin's bullet is aimed.

And that is obviously bogus. Every assassin's bullet has its target (let's suppose), without there being a single target shared by them all. Likewise, every practice may aim at some valuable end or other without there being a single good which encompasses them all.

Drawing out some logical structure, Aristotle's argument appears to rely on the inference form

Every F is R to some G
So: There is some G such that every F is R to it.

and then **C′** is a counterexample to the reliability of that form of inference.

(b) Did Aristotle really use an argument relying on that disastrous form of inference? Or can he wriggle off the hook?

We have ripped the quotation from Aristotle out of all context, so there is room for debate about what his intended argument really is. We can't enter into such debates here. Still, anyone who wants to defend Aristotle must at least meet the challenge of saying why the supposed counterexample fails to make the case, and must tell us what principle of inference really is being used in his intended argument here. If Aristotle isn't relying on the fallacious inference-type as displayed, then what *is* he up to?

The same goes, to repeat, for other challenges by the counterexample method. Given an argument in a text apparently relying on some identified pattern of inference, finding a counterexample to the deductive reliability of that pattern delivers a potentially fatal blow. If you want to defend the author, then your only hope is to show that the text has been misunderstood and/or the relevant inference step mis-identified.

(c) Expressions like 'every' and 'some' are standardly termed *quantifiers* by logicians (see Chapter **??**): what we have just noted is that we can not always shift around the order of quantifiers in a proposition.

It is perhaps worth briskly noting in passing that a number of naive arguments for the existence of God commit the same apparently tempting *quantifier-shift fallacy*, using the same bad form of inference. Consider, for example,

D (1) Every ecological system has an intelligent designer,
So (2) There is some intelligent designer (God) who designed every ecological system.

The premiss (1) was exploded by Darwin; but forget that. The point we want to make now is that even if you grant (1), that does not establish (2). There may, consistently with the premiss, be no one Master Designer of all the ecosystems – rather each system could be produced by a different designer. (No respectable philosopher of religion uses the rotten argument **D** as it stands; but you need to appreciate the fallacy in the naive argument here to see why serious versions of design arguments for God have to try a great deal harder!)

5.4 Summary

The counterexample strategy for proving invalidity is applicable to a target argument when (1) we can find a pattern of inference that the argument is depending on, and (2) we can show that the pattern is not a reliable one by finding a counterexample to its reliability, i.e. find an argument exemplifying this pattern which has (or evidently could have) true premisses and a false conclusion.

This method is familiar in the everyday evaluation of arguments as the 'But you might as well argue . . . ' gambit.

We used the counterexample strategy to show, in particular, that a certain (tempting?) quantifier-shift fallacy is indeed a fallacy.

Exercises 5

Some of the following arguments are invalid. Which? Why? In particular, use the counterexample strategy to prove invalid the ones that are.

(1) Many ordinary people are corrupt, and politicians are ordinary people. So, some politicians are corrupt.

(2) Many great pianists admire Glenn Gould. Few, if any, unmusical people admire Glenn Gould. So few, if any, great pianists are unmusical.

(3) Everyone who admires Bach loves the Goldberg Variations; some who admire Chopin do not love the Goldberg Variations; so some admirers of Chopin do not admire Bach.

(4) Some nerds are trainspotters. All trainspotters wear parkas. Some who wear parkas wear baseball caps too. So some trainspotters wear wear baseball caps.

(5) Anyone who is good at logic is good at assessing philosophical arguments. Anyone who is mathematically competent is good at logic. Anyone who is good at assessing philosophical arguments admires Bertrand Russell. Hence no-one who admires Bertrand Russell lacks mathematical competence.

(6) Everyone who is not a lunatic can do logic. No lunatics are fit to serve on a jury. None of your cousins can do logic. Therefore none of your cousins is fit to serve on a jury.

(7) Most logicians are philosophers; few philosophers are unwise; so at least some logicians are wise.

(8) All logicians are rational; no existentialists are logicians; so if Sartre is an existentialist, he isn't rational.

(9) If Sartre is an existentialist, he isn't a logician. If Sartre isn't a logician, he isn't good at reasoning. So if Sartre is good at reasoning, he isn't an existentialist.

6 Logical validity

We now introduce a narrower notion of validity, namely (so-called) *logical* validity. All logically valid arguments are deductively valid; but not vice versa.

Our discussion here will continue to be pretty informal. But it will be a useful stepping stone along the way to defining some more formal, precise, notions of particular kinds of logical validity that will in fact be our topic in the rest of this book.

6.1 Topic neutrality

Start by considering the following one-premiss arguments: are they deductively valid?

A (1) Jill is a sister.
So (2) Jill is female.

B (1) Jack has a first cousin.
So (2) At least one of Jack's parents is not an only child.

C (1) Jack is a bachelor.
So (2) Jack is unmarried.

In any possible situation in which Jill is a sister, she must be a female sibling, and hence be female. In any situation in which Jack has a first cousin (where a first cousin is a child of one of his aunts or uncles) he must have or have had an aunt or uncle, so his parents cannot both have been only children. Necessarily, if Jack is a bachelor, he is unmarried. So these arguments all involve deductively valid inference steps, according to our characterization of validity in §2.1.

In each of these cases, we can again abstract from some of the details, and see the argument as an instance of a schema all of whose other instances are valid too. For plainly, our arguments are not made valid by something peculiar to Jill or Jack; rather, what matters are the concepts of a sister, a first cousin, a bachelor. Therefore any inference of the following kinds is valid:

n is a sister
So: n is female

n has a first cousin
So: n's parents are not both only children

n is a bachelor
So: n is unmarried.

Here, then, we have some further examples of reliable principles of inference. But in an obvious sense, these principles are less abstract than those we have highlighted previously.

Let's say, as a (rough, first-shot) definition:

> Vocabulary like 'all' and 'some', 'and' and 'or', 'not' and 'if' etc. (plus the likes of 'is', 'are') – vocabulary which doesn't refer to specific things or properties, but which is useful in discussing any topic – is *topic neutral*.

The schemas we used in earlier chapters to display some inferential patterns involved only schematic letters and topic-neutral vocabulary. By contrast, the schemas just introduced involve concepts which belong to more specific areas of interest – these first examples in fact all concern familial relations. The new patterns of inference are none the worse for that; their instances are perfectly good valid inferences by our informal definition in §2.1. But their validity does not rely merely on the distribution of topic-neutral vocabulary in the premisses and conclusions.

Let's have a few more examples:

D (1) Kermit is green.
So (2) Kermit is coloured.

E (1) Cambridge is north of Oxford.
So (2) Oxford is south of Cambridge.

F (1) Siena is in Tuscany,
(2) Tuscany is in Italy,
So (3) Siena is in Italy.

The inference steps here are also all valid. We can again abstract to get shareable principles of inference that can recur in other arguments. But again, the schemas revealing these shareable principles will involve concepts which aren't topic-neutral (for example, the related concepts of *being green* and *being coloured*).

6.2 Logical validity, at last

(a) Having noted that there are arguments whose deductive validity depends on features of special-interest concepts such as those to do with family connections or being coloured etc., we will now largely put them aside. From here on, we will be concentrating on the kinds of valid arguments illustrated in our earlier examples – i.e. on arguments whose load-bearing principles of inference can be laid out using only topic-neutral notions. Such arguments involve patterns of reasoning that can be used when talking about any subject matter at all. And these universally-applicable patterns of reasoning are the special concern of logic as a discipline (yet another idea that goes back to Aristotle).

(b) It is useful to have some terminology to mark off these core cases of valid arguments which rely on topic-neutral principles of inference. Acknowledging the traditional focus of logic, we will say:

An inference step is *logically valid* if and only if it is both deductively valid (i.e. is necessarily truth-preserving) and the way in which topic-neutral notions occur in the premiss and conclusion is enough to ensure its validity.

Some premisses *logically entail* a certain conclusion if the inference from premisses to the conclusion is logically valid.

Here is an alternative way of defining logical validity. Let a *purely logical schema* be one involving only schematic letters and topic neutral vocabulary. (For convenience, we will allow limiting cases such as '*A, so A*', which is a schema built just from schematic letters. We will also allow limiting cases like '*There is something; so there is not nothing*' which lacks symbols and only features topic neutral vocabulary; this can count as a purely logical schema whose only instance is itself.) Then we can equivalently say: a logically valid inference is an instance of a purely logical schema all of whose instances are necessarily truth-preserving.

In this usage, the inferences in arguments **A** to **F** in this chapter, although deductively valid in our original sense, do not count as *logically* valid. Contrast our bold-labelled examples of valid arguments in earlier chapters – apart from one possible exception to which we return in a moment, those *are* logically valid.

(c) How can we show that an unobviously valid inference is logically valid in the sense explained? By a proof, of course, where the pattern of inference used at each step of the proof can be displayed using a purely logical schema. See for example our fully annotated proof labelled **D′** in §4.2: by inspection, this not only shows that inference from the initial premisses to the final conclusion is deductively valid but also that it is, more specifically, logically valid.

How can we demonstrate that an argument is *not* logically valid? If we can show that the argument is not deductively valid (by a counterexample, say), then that settles the question. But this point won't help us assess e.g. the Kermit argument which *is* deductively valid. However, in the case of that argument, we can note that the most structure we can expose with a purely logical schema is *n is F, So n is G*; and *that* is obviously not a reliable pattern of inference! So the Kermit argument is indeed not *logically* valid. Similar reasoning can be used in many other cases.

(d) Like our earlier definition of 'deductive validity' in §2.1, our new definition of 'logical validity' is, in one version or another, fairly standard.

Or rather, the concepts defined are standard; but note that the labels attached to these concepts by different writers can in fact vary. Moreover, those writers who foreground just *one* of these two concepts will tend to call that one, whichever it is, simply 'validity' (unqualified). So care is needed when comparing different textbook treatments.

6.3 Logical necessity

We say that an inference is logically valid if it is necessarily truth-preserving in virtue of how topic-neutral logical notions feature in its premisses and conclusion. Similarly:

A proposition is *a logically necessary truth* (or is simply *logically necessary*) if it is necessarily true in virtue of how topic-neutral logical notions feature in it.

So compare *Whatever is green is coloured* with *Whatever is both green and square is green*. The first proposition is true of every possible situation, so is necessarily true in the sense of §2.1(d). But it is not *logically* necessary in our sense, for its truth depends on the internal connection between being green and being coloured. But the second proposition is an example of the schematic pattern 'Whatever is F and G is F', any of whose instances has to be true just in virtue of the meanings of the topic-neutral logical notions 'whatever' 'is' and 'and'. For that reason, the second proposition is said to be *logically* necessary.

Generalizing, a logically necessary truth will be an instance of a type of proposition which can be specified by a purely logical schema using just schematic letters and topic-natural vocabulary, where all of the instances of that schema are necessarily true.

As noted back in §2.1, the notions of deductive validity and necessary truth are tied tightly together. Thus, the single premiss inference *A, so C* is valid if and only if it is necessarily true that *if A, then C*. We should now remark that, exactly similarly, the notions of logical validity and logical necessity also fit tightly together: *A, so C* is logically valid if and only if it is logically necessary that *if A, then C*.

6.4 The boundaries of logical validity?

We can say, in a summary slogan, that an argument is logically valid if it is valid in virtue of its topic-neutral form.

But how secure is our grasp on the notion of topic-neutrality here? We started making a list of topic-neutral vocabulary like 'all' and 'some', 'and' and 'or', etc.; but exactly how do we continue the list, and where do we stop? It in fact isn't entirely obvious what is to count as suitably topic-neutral vocabulary (without going into details, there are various stories on the market, but no agreed one). Hence it is not obvious what counts as a purely logical, topic-neutral, schema in our sense. Hence it is not obvious either just what will count as a *logically* valid inference, valid in virtue of its form as captured by a purely logical schema.

Here's a simple case to think about. Take the argument

G (1) Bill is taller than Chelsea,
(2) Chelsea is taller than Hillary,
So (3) Bill is taller than Hillary.

This is valid, and our first attempt at exposing a relevant general pattern of inference here might be

m is a taller then n
n is a taller then o
So: m is a taller then o.

And this isn't a purely logical, topic-neutral, schema.

But hold on! Can't we expose some more abstract inferential structure here? – like this, perhaps:

m is F-er than n
n is F-er than o
So: m is F-er than o.

Here *is F-er than* is the comparative of *F* (in other words, it is equivalent to *is more F than*). And *this* pattern of valid inference using the comparative-forming construction arguably *is* topic-neutral and so purely logical.

But hold on again! We in fact can only talk of one thing being more *F* than another thing, when *F* stands in for the kind of property that comes in degrees, as in *tall*, *wise*, *heavy*, *dark*, *flat*, etc. But we can't take comparatives of other properties like *even* and *prime* (of numbers), *dead* (of people), *valid* (of arguments), and so on. So arguably the comparative construction *is . . . -er than* is not an entirely topic-neutral device after all: it depends on what we apply it to whether it makes sense.

So is our last schema fully topic neutral or not? We can't pursue this issue further here; we will have to leave the question hanging, and so leave it unsettled whether **G** counts as *logically* valid. We will similarly leave it unsettled whether the 'Everyone loves a lover' argument **B** in §4.1 counts as logically valid (that will depend on how you think of the link between 'is a lover' and 'loves someone').

Fortunately, that's no problem. We don't need to worry about just how much the notion of purely logical validity covers. That's because our real concern will not be with some putative *general* notion of logical validity but rather with a number of important, very sharply defined, *particular* types of case. In other words, we will be considering various ways in which validity can depend on particular rather limited selections of quite uncontroversially topic-neutral vocabulary like 'all' and 'some', 'and' and 'or', etc . – so at least *these* cases will be absolutely clear examples of logical validity. Likewise, we will be focussing on clear, central, cases of logical necessity. Much more on this in due course, starting in real earnest in Chapters **??** and **??**.

6.5 Definitions of validity as rational reconstructions

Let's now put aside the issue of the exact scope of the idea of *logical* validity. That apart, how well does our classical definition of logical validity in §6.2 tally with our initial hunches about what makes for a completely watertight, absolutely compelling, inference step that depends on simple logical notions?

There are a number of potential issues here, but let's focus on just one for now. In fact, exactly the same worry already arises concerning our earlier definition of the wider notion of deductive validity in §2.1; but it would have been far too distracting to mention it there, right at the very beginning.

Consider, then, the following argument:

H Jack is married. Jack is not married. So the world will end tomorrow!

Plainly, just in virtue of the role of 'not' here, there is no possible situation in which the premisses of this inference are true together (being married and being not married rule each other out). Hence there is no possible situation in which the premisses of this inference are true together and the conclusion is false. Hence, applying our definition, **H** counts as a logically valid inference (as is any inference of the form *A, not-A, so C*, for any propositions *A* and *C*). Which might well seem a *very* unwelcome verdict.

One possible initial response runs along the following lines:

> Recall Aristotle's natural definition of a correct deduction as one whose conclusion "results of necessity" from the premisses. And indeed, surely the conclusion of any compelling deduction should have *something* to do with the premisses: so how could some claim about the end of the world really follow from propositions about Jack's marital status? Argument **H** commits a fallacy of irrelevance.
>
> Our definition of deductive validity in §2.1, and our derived definition of the narrower notion of logical validity in §6.2, therefore both overshoot – they count too many inferences as cogent. So, the definitions need to be tightened up by introducing some kind of relevance-requirement in order to rule out such daft examples as the apocalyptic **H** counting as 'valid'. Back to the drawing board!

The trouble is that when we get back to the drawing board, we find it is very difficult to respect the hunch that **H** is not a cogent inference without offending against *other*, equally basic, intuitions (see §**??**). And this prompts a different response to that initially unwelcome implication of our definition of validity:

> Arguably, *no* crisp definition of deductive validity (or of the special case of deductive-validity-due-to-topic-neutral-vocabulary) can be made to fit all our untutored hunches about what is and isn't an absolutely compelling inference. The aim of definitions of validity, then, is to give tidy 'rational reconstructions' of our informal concepts which smoothly capture the most important, most central, cases. Our informal definitions of the notions of deductive and logical validity do this neatly and naturally.
>
> True, we get odd results like the verdict on **H**. But this is a small price to pay in tidying up the notion of validity. After all, **H**'s premisses can never be true together; so we can't ever *use* this type of officially valid inference to establish the irrelevant conclusion as true.Our definition therefore sanctions only an entirely harmless extension of the intuitive idea of an absolutely compelling inference. Similarly (we hope!) we can live with a few other initially odd-looking implications of our official definitions.
>
> It will be similar for the technical notions of validity that we later introduce (like Chapter **??**'s 'tautological validity'). These again have the consequence that arguments like **H** are valid, but the definitions are defensible as rational reconstructions of our intuitive ideas which neatly characterize some special types of validity. More on this in due course.

This second line will be our response (and it is very much the majority response among modern logicians).

It is worth emphasizing that fixing an official definition for some logical notion on the basis of a sort of cost-benefit assessment is rather typical; we will meet other examples later. The best logical theory isn't handed down, once and for all, on tablets of stone. We repeatedly have to balance, say, a certain natural simplicity or smoothness of theory over here against a certain artificiality over there. And the best choices for definitions, i.e. the best choices for rational reconstructions of informally messy ideas, can depend on our particular theoretical priorities in a given context. This is a theme we will need to return to.

6.6 Summary

Vocabulary like 'all' and 'some', 'and' and 'or', etc. (plus the likes of 'is', 'are'), not about particular things or properties, but useful in discussing any topic, is said to be topic neutral.

A purely logical schema is one involving only schematic letters and topic-neutral vocabulary.

A logically valid inference is an instance of a purely logical schema all of whose instances are necessarily truth-preserving.

A logically necessary proposition is an instance of a purely logical schema all of whose instances are necessarily true.

It isn't obvious what counts as topic-neutral vocabulary, so not clear either where to draw the boundaries of logical validity/logical necessity.

Exercises 6

Which of the following arguments are deductively valid? Which are logically valid?

(1) Mr Hyde is none other than Dr Jekyll. Mr Hyde is six foot tall. Hence Dr Jekyll is six foot tall.

(2) There are six oranges in the bowl. Hence there is an even number of oranges in the bowl.

(3) There are infinitely many stars. Every star is a gaseous body. Hence there are infinitely many gaseous bodies.

(4) There is liquid H_2O in the cup. Hence there is water in the cup.

(5) The Battle of Hastings happened before the Battle of Waterloo. The Battle of Marathon happened before the Battle of Hastings. Hence The Battle of Marathon happened before the Battle of Waterloo.

(6) There are more people than there are hairs on any one person's head. Hence at least two people have the same number of hairs on their head.

7 Propositions and forms

One last topic before we get down to formal business.

Propositions – the ingredients of arguments – are, we said, the sort of thing that can be true and false (as opposed to commands, questions, etc.). But what is their nature? In this chapter we briefly outline two common types of view. We also explain why (for our purposes, at any rate) we happily don't need to decide between them. Nor need we resolve a related question about the nature of *forms* of everyday inference.

We begin, however, with a definition and two distinctions.

7.1 A definition: Fregean sense

We will say that the *sense* of a sentence – a sentence apt for propositional use in stating a premiss or concluision – is that aspect or ingredient of its meaning that is relevant to questions of truth or falsity. This use of 'sense' is due to the great nineteenth-century German logician Gottlob Frege (translating his *Sinn*).

So the Fregean sense of a sentence fixes the condition under which it is true. Distinguish this from the 'colouring' or 'flavour' or *tone* which the choice of different words might give a sentence. Whether I refer to a particular woman as 'Elizabeth' or 'Lizzie' may reflect my closeness (or my lack of respect). And, to use an example of Frege's, whether I refer to her mount as a 'horse' or 'steed' or 'nag' or 'gee-gee' may reflect my estimate of the beast (or view of Elizabeth's riding). But these differences of tone don't affect the truth or falsity of e.g. 'Elizabeth's/Lizzie's horse/steed/nag/gee-gee is black' – the various permutations have the same truth-relevant sense, and hence are true in the same worldly situations. Or at least, so goes a very natural story. And we don't have to buy any particular theory about sense (such as Frege's own) to acknowledge that we need to make *some* such distinction between core factual meaning and its embellishments.

Now, as far as logic is concerned, questions of colouring or tone aren't going to matter – those aspects of the overall meaning of claims don't affect their logical properties. Because what matters for validity is preservation of unvarnished truth. So it is the *sense* of sentences that we will care about.

7.2 A distinction: types vs tokens

We now remark that we talk about sentences in two ways – the distinction must surely be an ancient one, but our terminology is due the nineteenth-century American philosopher

C. S. Peirce. It is best introduced via an example.

Suppose then that you and I take a piece of paper each, and boldly write 'Logic is fun!' a few times in the centre. So we produce a number of different physical inscriptions – perhaps yours are rather large and in blue ink, mine are smaller and in black pencil. Now we key the same encouraging motto into our laptops, and print out the results: we get some more physical inscriptions, first some formed from pixels on our screens and then some formed from printer ink.

How many sentences, then? We can say: many, some in ink, some in pixels, etc. Equally, we can say: one sentence, multiply instantiated. Evidently we must distinguish the many different instances of the sentence – physically constituted in various ways, of different sizes, lasting for different lengths of time, etc. – from the linguistic pattern which they are all instances of.

Or to put that in Peirce's now-standard terminology: we have to distinguish the various physical sentence *tokens* from the sentence *type* which they are tokens of. (We can of course similarly distinguish word tokens from word types, book tokens – printed copies – from book types, and so on.)

What exactly makes a physical sentence a token of a particular type? – now, *there* is a tough question! And what exactly is the metaphysical status of types? – another tough question we can't worry about now. But it is agreed more or less on all sides that we need some kind of type/token distinction, however it is to be elaborated.

7.3 A distinction: propositions vs assertions

Go back for a moment to thinking about indirect arguments. Take our first example of a reductio ad absurdum argument in §4.5, which (renumbering) involved the step

A	(1)	Nobody is despised who can manage a crocodile.	(premiss)
	(2)	Babies *can* manage a crocodile.	(supposition)
	(3)	Babies are not despised.	(from 1, 2)

The indented subproof, recall, then continued until the supposition was 'discharged'. Now compare this bit of reasoning with the simple argument

B	(1)	Nobody is despised who can manage a crocodile.	(premiss)
	(2)	Babies *can* manage a crocodile.	(premiss)
	(3)	Babies are not despised.	(from 1, 2)

There's an obvious sense in which the ingredients of these two bits of reasoning so far are just the same; the same propositions occur as inputs to the same inference move, and the same proposition is derived. The difference is that, in the second case – at least when the argument is seriously propounded – the proposition at B(2) is *asserted* as a premiss; while in the first case, the same proposition at A(2) is merely temporarily *supposed* for the sake of argument.

Propositions, then, may occur both asserted and unasserted in arguments. Hence we must distinguish propositions as ingredients of arguments from assertions (a point famously emphasized by Frege).

7.4 Propositions as sentences, naively

The idea of sense, the type/token distinction, and the assertion/supposition distinction all deserve more careful discussion. But not here: we can't do everything at once! Instead, we return to the central philosophical issue of this chapter: what sort of things are the propositions that can feature as premisses and conclusions of everyday arguments?

(a) According to one view,

> Propositions, potential premisses and conclusions and the bearers of truth and falsity, are *declarative sentences* (meaning sentences like 'Jack kicks the ball' as opposed to interrogatives like 'Does Jack kick the ball?' or imperatives like 'Jack, kick the ball!').

Yet it seems quite implausible to identify premisses and conclusions with sentences in the sense of particular sentence *tokens*, particular physical inscriptions in ink, pencil or pixels. Consider this argument:

> All philosophers are eccentric. Jack is a philosopher. So Jack is eccentric.

We've met this argument before, haven't we? It's the same argument that we first encountered in §1.2. But there are of course different physical inscriptions here and there. So same premisses and conclusion, different sentence tokens.

Generalizing, we want to say that the same proposition can occur as a premiss on various pages in the many different printed copies of this book, and in the e-copies too. So, the view that propositions are sentences is surely better understood as holding that propositions are declarative sentence *types*, types which can have many scattered instances occurring on different pages in different books.

(b) Take, however, an ambiguous sentence-type like the proverbial wartime headline 'Eighth Army push bottles up Germans'; are the Germans being bottled up by a push, or are the Eighth Army pushing bottles? The sentence can say two different things, one of which may be true and the other false: so, same sentence, different truth-evaluable messages conveyed, different possible ingredients of arguments. Or for a more banal and rather more tasteful example, consider 'He loves her', occurring in varying contexts: again, the same declarative sentence can surely express different premisses/conclusions in different contexts.

Conversely, different sentences can express the same premisses/conclusions. For a trivial example, the arguments 'Jack is both tall and slim. So, Jack is tall.' and 'Jack is tall and is slim. So, he is tall.' involve distinct sentences; but it seems extremely odd to say that the stylistic variation makes them different arguments. Another case (which alludes to an example of Frege's): consider the arguments 'The Greeks defeated the Persians at Palatea. So the Persians were defeated at least once.' and 'The Persians were defeated by the Greeks at Palatea. So the Persians were defeated at least once.' Aren't these the same argument, with the active and passive versions of the first premiss again just stylistic variants – the two sentences having the same sense, expressing the same thought?

Those first examples only involve modestly massaging the sentences in premisses and conclusions. Much more radically, though, can't the same argument appear in this book and in its Italian translation, with the same premisses and conclusions now expressed by entirely different sentences?

In short, then, it seems that the same argument, with the same propositions as premisses and conclusion, can be presented using different sentence types, so long as the sentences continue to have the same sense so that the messages expressed stay the same. Therefore propositions aren't just sentence types either.

7.5 Propositions as truth-relevant contents

Considerations like those last ones push us towards a second view, along the following lines:

> Propositions, potential premisses and conclusions and the bearers of truth and falsity, are strictly speaking not declarative sentences but the messages that the declarative sentences can be used, in context, to express.

Some would talk of 'thoughts' rather than 'messages' as what are expressed by declarative sentences – meaning possible thought-contents as opposed to acts of thinking. Some talk just of 'contents'. And there is usually some further elaboration – it is the truth-relevant content (i.e. what gets determined by the *sense* of a sentence, ignoring questions of tone) that matters here. So a proposition, on this second story, is identified not by linguistic shape but by the conditions under which it is true.

(An aside. We have been using the word 'proposition' in a non-committal theory-neutral way – so making it an open question what propositions are. Confusingly, many instead use 'proposition' specifically to refer to these truth-relevant thought-contents expressed by sentences. Though just to complicate things further, medievals and modern writers influenced by them use 'proposition' in exactly the opposite way, specifically to refer to declarative sentences themselves. Sorry about that – you just have to be alert when reading other authors!)

Whatever the terminology, the trouble with our second view about the bearers of truth and falsity is that it tells us what premisses and conclusions are not – to repeat, on this view they aren't sentences. However, their positive nature is left rather mysterious. In fact it is quite unclear what kind of theory to offer about messages or thought-contents, and no attempt to give one commands very wide support.

7.6 Why we can be indecisive

So maybe we should after all try to rescue the less puzzling view that the bearers of truth are sentences by making some suitable adjustments. Consider, then, the revised suggestion that propositions are *fully interpreted* declarative sentences (i.e. they are disambiguated sentences parsed as having a determinate Fregean sense, with context supplying the references of pronouns etc.). Then we may equate e.g. 'Jack is tall' and 'he is tall' when the pronoun refers to Jack; and maybe we can allow other grammatical massaging, e.g. equating active and passive versions of the same sentence. On the other hand, we could perhaps bite the bullet and insist that the premisses and conclusions of an argument in this book and the corresponding premisses and conclusions in its Italian translation are not strictly speaking the same after all, as the sentences are so different – rather, there are two distinct arguments which are more or less smooth translations of each other (though translation is famously not an all-or-nothing business, but comes in degrees of acceptabiity).

How, though, are we to nail down some version of this revised suggestion, based on the idea that propositions are fully interpreted sentences?

Another very good question: and we will just have to leave it hanging. Having flagged up that there are different sorts of views about how we should think of premisses and conclusions in everyday arguments, we aren't going to try to answer the question of which approach is right and how best to develop it.

Sitting on the fence about the nature of ordinary propositions, i.e. about the nature of premisses and conclusions in everyday arguments, may sound simply irresponsible; how can logicians possibly leave unresolved such a very basic question about arguments as what they are made of? For our purposes in this book, however, it turns out that we happily won't need to adjudicate this tricky issue in 'philosophical logic'.

Why so? Because – spoiler alert! – our key technique for assessing everyday arguments will involve rendering them first into disciplined formalized languages, and then concentrating our logical efforts on assessing arguments once neatly formalized. *But sentences of our formalized languages will by design be free of ambiguities and context dependence* (they come, so to speak, already 'fully interpreted', stipulated to have a determinate sense). Formalized sentence types and the messages they convey will therefore be neatly aligned. Hence, as we will see, it will do no harm to take arguments framed in *formalized* languages to be made up of their formal sentences in a straightforward way. In other words, in the special context of dealing with formalized arguments, our main topic in what follows, the propositions making up arguments *can* – at least for our purposes – be unproblematically identified with sentence-types plain and simple.

7.7 Forms of inference again

(a) In discussing patterns or forms of inference, we have so far been very relaxed and informal – but hopefully clear enough for our introductory purposes in these early chapters. However, we have just asked how we should think of the premisses and conclusions of everyday arguments: are they sentences or are they the messages expressed by sentences (i.e. thought-contents)? We can now raise a related issue: when we talk about forms of inference in everyday arguments, are we referring to patterns to be found at the surface level of the sentences used to state the argument, or should we primarily be thinking of patterns at the level of messages expressed (whatever exactly that means)?

In fact, we have been cheerfully casual about this. On the one hand, we have discussed patterns of inference to be found in arguments couched in English by using schemas in which we are to systematically substitute English expressions for the schematic letters – allowing for some grammatical tidying. Which chimes with the policy of thinking of the ingredients of arguments as sentences, and with thinking of patterns of inference as patterns to be found on the surface, at sentence level.

On the other hand, when we laboriously stated in words the principles underlying our first couple of examples of reliable inference schemas (in §1.5 and §3.1), we found ourselves talking about what premisses and conclusions *say*, and this looks to be at the level of messages expressed. It is very natural to slide into this way of speaking. After all, think of Aristotle writing about arguments in Greek or medievals writing in Latin or Frege writing in German. Can't those great dead logicians, who obviously were considering arguments couched in very different sentences, still be discussing the very same forms of inference that occur in our arguments couched in English?

(b) Leaving aside arguments in other languages, the issue already arises within a single language. For example, should we think of the following four arguments as sharing the same pattern of inference?

	All dogs have four legs.		Every dog has four legs.
	Fido is a dog.		Fido is a dog.
So,	Fido has four legs.	So,	Fido has four legs.
	Any dog has four legs.		Each dog has four legs.
	Fido is a dog.		Fido is a dog.
So,	Fido has four legs.	So,	Fido has four legs.

Thinking at the level of sentences, these arguments exemplify different forms of inference, because they involve the distinct logical words 'all', 'every', 'any' and 'each'. However we might well be inclined to suppose that the differences in these cases are only superficial and that the underlying inferential structure is in some sense really the same. The respective first premisses of these arguments are just stylistically different ways of expressing the very same thought-content; and then the rest of the arguments are the same. It is tempting, then, to suppose that although these various versions may differ in surface sentential form, they in *some* way share the same underlying 'logical form' – so, thinking at the level of the messages expressed by the various sentences, the inferences are all instances of a single form. We might even recruit the familiar schema '*All F are G, n is F, so n is G*' to represent this shared underlying form: schemas are very often used in this way in philosophical writing, i.e. are treated as sitting quite loose to the surface form of arguments.

(c) So which should it be, when thinking about the forms of inference in everyday arguments? Should we be looking for patterns at sentence-level, or for underlying patterns in the messages expressed? As we have just seen, the second line can be tempting. But do we really understand what it comes to – for what kind of theory of the nature of messages would be needed for it to make sense? There is no widely agreed account of propositions-as-messages that we can appeal to for help here.

Fortunately, at an introductory level, it doesn't matter much whether we say e.g. that Aristotle was discussing the very same forms of argument as we might use, or whether we instead say that he was considering the Greek equivalents of our forms of argument. And much more importantly – another spoiler alert! – when we start to consider arguments regimented into formalized languages, things become very clean and simple. As we said, formal arguments can be thought of as made up of formal sentences; and then the inferential forms of such arguments can be understood quite unproblematically as formal patterns at the sentential level. More about this soon.

7.8 Summary

There are two main types of view on the nature of the propositions in ordinary arguments: they are sentences (perhaps fully interpreted sentence-types) or alternatively are the messages or thought-contents that sentences can be used to express.

We need not adjudicate: our focus from now on will be on propositions in formalized arguments, and these can unproblematically be treated as sentences.

Likewise, there are differing views about what forms or patterns of inference are patterns in – sentences or messages?

Again we do not need to adjudicate: our focus from now on will be on forms of inference in formalized arguments, and these can unproblematically be treated as patterns in the surface form of formal sentences.

Interlude: From informal to formal logic

What have we done so far? In bare headlines,

> We have explored, at least in an introductory way, the (classical) notion of a valid inference-step, and the corresponding notions of a deductively valid/sound (one-step) argument.
>
> We have seen how to distinguish deductive validity from other virtues that an argument might have (like being a highly reliable inductive argument).
>
> We have noted how different arguments can share the same form of inference. And we have seen how to exploit this fact in the counterexample method for demonstrating invalidity.
>
> We have seen some simple examples of direct multi-step proofs, where we show that a conclusion really can be validly inferred from certain premisses by filling in the gap between premisses and conclusion with evidently valid intermediate inference steps. We in addition briefly looked at one kind of indirect method of proof, reductio ad absurdum.
>
> We have also met the narrower notion of logical validity – where an inference is valid in this narrower sense if it is deductively validity in virtue of the way that topic-neutral notions feature in the premisses and conclusion.

Along the way, we have had to quietly skate past a number of issues, leaving more needing to be said. But hopefully, you will have gained at least a rough-and-ready preliminary understanding of some key logical concepts.

So what next? One good option would be to spend more time on techniques for teasing out the inferences involved in passages of extended prose argumentation, to develop further methods of informal argument analysis, explore how to reason with a range of key logical notions like ‘if’ and ‘all’, and then catalogue a variety of common ways in which everyday arguments can go wrong. This kind of study in *informal logic* (as it is often called) can be a highly profitable exercise.

But our focus will be different. Instead of going for breadth of coverage, we aim for depth; we will develop systematic and rigorous treatments for some highly important but limited classes of arguments. Moreover – and this is a crucial move – the arguments we focus on will be formalized arguments, framed in purpose-designed formalized languages. It will quickly become clear why we are making this move, starting in the next chapter. For the moment, I just claim that the resulting theory of *first-order quantification theory* – the main branch of formal logic which we are eventually introducing in this book – is one of the great intellectual achievements of formally minded philosophers and of philosophically minded mathematicians. It is beautiful in itself, and it opens the door onto a very rich and fascinating field. (In this introductory book, there will only be very occasional glimpses further through that door; but we will at least get to the threshold.)

We will take things slowly, however. Quantification theory explores the logic of arguments involving expressions of generality ('all', 'some', and 'none' etc.) and also explains how these expressions interact with the so-called connectives ('and', 'or, 'if', 'not'). But before turning to the full theory, we are going to be spending a *lot* of time and effort on the fragment of this logical system which just deals with the connectives, i.e. *propositional logic*. Why is this worth doing at length? Yes, propositional logic does have some intrinsic interest. But we explore it in detail principally because a whole range of basic ideas and strategies of formal logic can be most accessibly introduced and initially explored in the context of this cut-down theory. Which makes for a relatively painless introduction to our subject, and will *very* considerably ease the path into full quantification theory.

So let's get straight to work!