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*.

But 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, we also deploy reasoned arguments in the course of 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, call it argument **A**. Suppose you hold

(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 can now be evaluated along two quite independent dimensions:

> First, we can ask whether argument **A**'s *premisses* (1) and (2) are true. In other words, are the 'inputs' to your inference step 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 you from **A**'s 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 – the logician wants to know when an argument's premisses, supposing that we accept them, do indeed provide compelling grounds for also accepting its conclusion. It is in this sense that logic is concerned with the 'internal cogency' of our reasoning.

1.3 Deduction vs. induction

(a) We said that argument **A**'s inference step is absolutely compelling: 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 (to illustrate that not all good reasoning is of this type). 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 rather 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 the 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 this inference step 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 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 have to trust our lives, day in, day out. 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 deductively 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 people are Catholics.
(2) Most Catholics oppose abortion on demand.
So (3) At least some Irish people 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). Hence the inference step is not deductively valid.

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

F (1) Some philosophy students admire all logicians.
(2) No philosophy student admires anyone irrational.
So (3) No logician is irrational.

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 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: