1

What is logic?

The business of logic is the *systematic evaluation of arguments for internal cogency*. And the kind of internal cogency that will especially concern us is *deductive validity*. But these brief headlines leave everything to be explained.

- What do we mean here by 'argument'?
- What do we mean by evaluation for 'internal cogency'?
- What do we mean, more particularly, by 'deductive validity'?
- What sorts of 'systematic' evaluation of arguments are possible?

This introductory chapter makes a start on answering these questions.

1.1 What is an argument?

By 'argument' we mean, roughly, a chain of reasoning in support of a certain 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, i.e. they are not yet giving any reasons in support of one claim or the other.

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

1.2 What sort of evaluation?

The business of logic, then, is the evaluation of stretches of reasoning. Let's take a very simple case (call it argument **A**).

Suppose you hold

(1) All philosophers are eccentric.

I then introduce you to Jack, who I tell you is a philosopher. So you come to believe

(2) Jack is a philosopher.

Putting these two thoughts together, you can obviously draw the conclusion

(3) Jack is eccentric.

This little bit of reasoning can now be evaluated along two quite independent dimensions.

- First, we can ask whether the *premisses* (1) and (2) are true: are the 'inputs' to your reasoning correct? (1) in fact looks very disputable. And maybe Jack's reflective skills are so limited that we'd want to dispute the truth of (2) as well.
- Second, we can ask about the quality of the *inference* from the premisses (1) and (2) to the *conclusion* (3). In this case, the movement of thought from the premisses (1) and (2) to the conclusion (3) is surely absolutely compelling. We have agreed that it may be open to question whether (1) and (2) are both true. However, *if* (1) and (2) are granted to be true (granted 'for the sake of argument', as we say), then (3) has got to be true too. There's just no way that (1) and (2) could be true and yet (3) false. Someone who asserted that Jack is a philosopher, and that all philosophers are eccentric, yet went on to *deny* that Jack is eccentric, would be implicitly contradicting himself.

In brief, it is one thing to consider whether an argument starts from true premisses; it is another thing entirely to consider whether it moves on by reliable inferential steps. To be sure, we normally want our arguments to pass muster on both counts. We normally want to start from true premisses *and* to reason by steps which will take us on to further truths. But it is important to emphasize that these are distinct aims.

The premisses (and conclusions) 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 matter of physics, then the physicist is the one who can help us about the truth of our premisses; and so on. The specific concern of logic, by contrast, is not the truth of initial premisses but *the way we argue from a given starting point*. Logic is not about whether our premisses are true, i.e. match up to the world, but about whether our inferences really do support our conclusions once the premisses are granted. It is in this sense that logic is concerned with the 'internal cogency' of our reasoning.

1.3 Deduction vs. induction

The one-step argument **A** is a splendid bit of reasoning; if the premisses are true, then the conclusion is *guaranteed* to be true too. Here's a similar case:

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

Who knows whether the premisses are true or not? But we can immediately see that necessarily, *if* the premisses are true, then the conclusion will be true too: the inferential move is absolutely watertight. If premisses B(1) and B(2) here are both true, then B(3) cannot conceivably fail to be true.

Now consider the following contrasting case. Wild philosophical scepticism apart, you are thoroughly confident that the cup of coffee you are drinking isn't going to kill you (and 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 something like

C (1) Cups of coffee that looked and tasted just fine haven't killed you in the past.

You also know that

(2) This present cup of coffee looks and tastes just fine.

These premisses will prompt you to conclude

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

And the inference that moves from the premisses C(1) and C(2) to the conclusion C(3) is, in the circumstances, surely perfectly reasonable: 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 crucial contrast with the earlier 'Jack' and 'Jill' examples – it *isn't* 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 generally been in the past don't guarantee that the trend will continue in the future.

Fortunately for you, C(3) no doubt *is* 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 deny C(3) might be saying something highly improbable given her premisses, but she won't actually be contradicting herself. We can make sense of the idea of C(1) and C(2) being true and yet C(3) false.

In his little book *The Problems of Philosophy*, Bertrand Russell has a nice example which presses home the same basic point. Here is the chicken in the farmyard. Every day of its life, the farmer's wife has come into the hen yard in the morning to feed it. Here she comes again, bucket in hand. The chicken believes that it is about to be fed, on the best possible evidence. But tough luck: today is the day when it gets its neck wrung. Again, what's happened in the past is a very good guide to what will happen next (what else can we rely on?): but reasoning from past to future isn't absolutely watertight.

In summary, then, there is a crucial 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 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) 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.

We'll need some terminology to mark this difference: so, as a first shot definition, we will say

> An inference step is *deductively valid* just if, given that its premisses are true, then its conclusion is absolutely guaranteed to be true as well.

Equivalently, when an inference is deductively valid, we'll say that the premisses *logically entail* the conclusion. So: arguments **A** and **B** each involve deductively valid inferential steps, but the coffee argument **C** doesn't. The latter argument instead involves reasoning from previously examined cases to a new case: this kind of extrapolation from the past to the future, or more generally from old cases to new cases, is standardly called *inductive*.

Note, 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 generally reliable reasoning to which we rightly trust our lives, day in, day out. The conclusion is highly likely to be true given the premisses. It is just that the inference in this case doesn't *guarantee* that the conclusion is true, even assuming the given premisses are true.

We'll say very little more about inductive inferences in this book. They are an important and difficult topic, but they are not *our* topic. Our concern is with arguments which aim to use deductively valid inferences, where the premisses are supposed to logically entail the conclusion. And our principal task will be to find various techniques for establishing whether a putatively valid inferential move really *is* deductively valid.

1.4 More examples

The 'Jack' and 'Jill' arguments were examples where the inferential moves are *obviously* deductively valid. Take the inference

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

That 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.

But questions of deductive validity needn't be *immediately* obvious, one way or the other. Suppose someone holds both of

E (1) Most Irish are Catholics.
(2) Most Catholics oppose abortion.

On those grounds he concludes

So (3) At least some Irish oppose abortion.

What are we to make of *this* argument? I don't know whether the premisses are true or not, and anyway that is hardly a matter for the logician (it is an empirical sociological matter to determine the distribution of religious affiliation among the Irish, and to find out how many Catholics support their church's official teaching about abortion). So let's leave aside the question of whether the premisses are in fact correct. What we *can* seek to determine here – from our armchairs, so to speak – is whether the inferential move is valid: *if* the premisses are true, then must the conclusion be true too?

Suppose the Irish are a tiny minority of the world total of Catholics. Then it could be that nearly all the other (non-Irish) Catholics oppose abortion, and hence most Catholics do, even though *none* of the Irish oppose abortion. In other words, E(1) and E(2) could be true, yet E(3) false. The truth of the premisses doesn't absolutely guarantee the truth of the conclusion, so the inference step can't be a deductively valid one.

Here's a last example:

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

Does this argument make a valid inferential move?

So far, all we can do to find the answer is to cast around experimentally, trying to work out somehow or other whether the truth of the premisses would guarantee the truth of the conclusion – though, as with the 'Irish' argument, a little thought should give the right answer (which is revealed in §3.1). Still, it would be good to be able to proceed more methodically and find some *general* techniques for evaluating arguments like this. Ideally, we would like general techniques that work *mechanically*, that can be applied to settle questions of validity as automatically as we can settle arithmetical questions by calculation. But are such techniques for systematically evaluating arguments ever available?

Later we will explore some general techniques – though we will also have occasion to note that there are intriguing limitations on what can be done mechanically. First, however, we should say a little more about what makes any kind of systematic approach possible.

1.5 The systematic evaluation of arguments

Consider again our first sample argument with its deductively valid inference step:

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

And 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 existentialists write bosh.
(2) Sartre is an existentialist.
So (3) Sartre writes bosh.

(We can evidently keep on going *ad nauseam*, churning out arguments to the same A-pattern!)

The four displayed arguments – and, of course, any others cast in the same mould – are all equally good in the sense of all involving valid inference steps. In each case the truth of the premisses would logically guarantee the truth of the conclusion. And plainly, it is no mere accident that the arguments are all internally cogent. The validity of the inference move in the first argument hasn't anything especially to do with being a philosopher or being eccentric. Likewise the validity of the inference move in the 'robin' argument isn't due to special facts about robins. The four arguments all work for the same reason: the shared pattern of inference is a reliable one.

We can describe the general principle here as follows.

> Any inference step which moves from a pair of premisses, one of which says that everything of a certain kind has a given property, and the other of which says that a particular individual is of the former kind, to a conclusion which says that that individual has the latter property, is deductively valid.

A moment's reflection shows that this principle is a correct one. But we have hardly presented it in the most perspicuous way. How much easier to put it instead like this:

> Any inference step of the following type
>
> *All F are G*
> *n is F*
> *So, n is G*
>
> is deductively valid.

Here the italic letters '*n*', '*F*', '*G*' are used to exhibit the skeletal pattern of an argument (with '*n*' holding the place for a name, '*F*' and '*G*' for predicates, i.e. expressions which attribute properties or pick out kinds of thing). How we flesh out the skeletal pattern or *schema* doesn't matter: any sensible way of substituting for the *schematic variables* (as we'll call the likes of '*n*', '*F*', '*G*') will yield another argument with a valid inference step of the same type.

We said at the outset, logic is concerned with the systematic study of validity. We now get a first glimpse of how systematicity might be achieved – by noting

that the same *patterns* or *forms* of inference can appear in many different particular arguments.

Some inference patterns like the one just shown are reliable. Others aren't. For example, consider this pattern of inference:

Most F are G
Most G are H
So, at least some F are H

This is the pattern of inference the 'Irish' argument was relying on, and we saw that it doesn't work.

There will be much more on the idea of patterns of inference in later chapters; but first, a quick summary.

1.6 Summary

- We can evaluate a piece of reasoning in two distinct ways. We can ask whether the premisses are actually true. And we can ask whether the truth of the premisses actually supports the truth of the conclusion. Logic is concerned with the second dimension of evaluation: i.e. the province of logic is the question whether a given argument internally hangs together, whether its inferential moves are cogent ones.
- We are setting aside inductive arguments (and other kinds of non-conclusive reasoning). We will be concentrating on arguments involving inferences that purport to be *deductively valid*. In other words, we are going to be concerned with the study of inferences that aim to *logically entail* their conclusions – aim to be, as it were, absolutely watertight. (Imagine an argument as a hydraulic device: if truth is poured in at the top, then we want it to trickle down to the bottom. The arguments with non-conclusive, possibly leaky, inference moves are ones where putting truth in at the top doesn't guarantee getting any truth out at the bottom; the deductively valid inferences are those where *truth in* guarantees *truth out*.)
- Arguments typically come in families whose members share good or bad types of inferential move; looking at such general patterns of inference will be a step towards making logic a systematic study.

Exercises 1

By 'conclusion' we do *not* simply mean what concludes, what is stated at the end of, a passage of reasoning. We mean what the reasoning aims to establish – and that might in fact be stated at the outset of an argument. 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 claim might well be given after the target conclusion has been stated.

Indicate the premisses and conclusions of the following arguments. Which of these arguments do you suppose involve deductively valid reasoning? Why? (We

haven't developed any techniques for you to use yet: 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. For 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 has ever overturned even a lead of 10 points with only a week to go to polling day.
8. 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.
9. No experienced person is incompetent. Jenkins is always blundering. No competent person is always blundering. Therefore Jenkins is inexperienced.
10. Many politicians take bribes. Most politicians have extra-marital affairs. So many people who take bribes have extra-marital affairs.
11. (Lewis Carroll) Babies cannot manage crocodiles. Because babies are illogical. But illogical persons are despised. And nobody is despised who can manage a crocodile.
12. (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.
13. '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.'
14. '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

The first chapter briefly introduced the idea of an inferential move being deductively valid. This chapter explores the idea of validity in a little more depth, and also emphasizes the centrality of deductive reasoning.

2.1 Validity and possibility

We said that an inference step is valid just if, given the input premisses are true, then the output conclusion is absolutely guaranteed to be true as well. But what do we mean here by an 'absolute guarantee'?

Previously, we gave a number of other, equally informal, explanations of the idea of validity in terms of what is 'conceivable' or what it would be 'self-contradictory' to assert. Here now is another definition of validity: and this time, it is a version of the standard, 'classical', textbook definition of validity.

> An inference step from given premisses to a particular conclusion is *(classically) valid* if and only if there is no possible situation in which the premisses would be true and the conclusion false.

We'll take this as our official definition from now on. But it is only as clear as the notion of a 'possible situation'. It needs to be stressed, then, that 'possible' here is meant in the widest sense – a sense related to those ideas of what is 'conceivable', or 'involves no contradiction' that we met before.

Consider, for example, this inference:

Premiss:	Jo jumped out of a twentieth floor window (without parachute, safety net, etc.) and fell freely and unimpeded onto a concrete pavement.
Conclusion:	Jo was injured.

And let's grant that there is no situation that can really obtain in the actual world, with the laws of physics as they are, in which the premiss would be true and the conclusion false. In the world as it is, falling unimpeded onto concrete from twenty floors up will always produce serious (surely fatal) injury. Does that make the inference from the given premiss to the conclusion deductively valid?

No. Although it isn't, let's agree, *physically* possible to jump in those circumstances without being injured, it remains possible in a weaker sense: it is '*logically* possible', if you like. We can coherently conceive of situations 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. There is no internal logical contradiction in that idea. *And all that is required for the deductive invalidity of an inference is that it is possible in this very weak sense for the premisses to be true and the conclusion false*. As we put it before, an inference is invalid if there is no internal contradiction in the notion of the truth of the premisses being combined with the falsity of the conclusion.

This very weak notion of logical possibility (any consistent fantasy counts as possible in this sense) goes with a correspondingly strong notion of logical *im*possibility. If something is *logically* impossible, like being circular and square at the same time, then it is absolutely ruled out: the very idea is inconsistent, incoherent, logically absurd. For a classically valid inference, it is impossible in this very strong sense for the premisses to be true and the conclusion false.

It will be useful to introduce a further definition:

> A set of propositions is *logically consistent* if it is logically possible for the propositions all to be true together. Likewise, a set of propositions is *logically inconsistent* if it is logically impossible for the propositions to be true together.

If there is no possible situation in which a bunch of premisses would be true and a certain conclusion false, then the premisses and the *denial* of the conclusion form an inconsistent set of propositions. So another way of characterizing the notion of validity is this: an inference is classically valid if the premisses taken together with the denial of the conclusion form an inconsistent set.

To be sure, we have still only gestured at the notions of logical possibility and logical impossibility that are in play here. As we go along, we will need to sharpen up our understanding of these ideas. Still, we've said enough to get us started.

2.2 What's the use of deduction?

It needs to be stressed again that deductively valid inferences are *not* the only acceptable ones. Inferring that Jo is injured from the fact she fell twenty stories onto concrete is perfectly reasonable. Inductive reasoning like this is very often good 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 – call him Sherlock. In the ideally satisfying detective story, Sherlock assembles a series of clues and then solves the crime by an *inference to the best explanation*. In other words, the detective arrives at an hypothesis that neatly accommodates all the strange events and bizarre happenings, an hypothesis that