

This argument has been used a good deal by British politicians in recent years, but is it a good argument? Certainly the reason given for the conclusion is *true*, but does the conclusion *follow* from it? Or could the reason be true and the conclusion false? It certainly doesn't follow from the fact that 'some people have done x' that 'everyone *could* do x': some people have run a mile in under four minutes but it doesn't follow that everyone could. Whether this analogous argument shows that our original argument is a bad one depends on whether it really is analogous – whether it exhibits the same logic. Does the original argument assume the principle that 'Some As are Bs so all As could be Bs' (clearly an unsound principle in general) or does it assume 'Some people have done x so everyone *could* do x' (clearly also an unsound principle) or is the argument specifically about unemployment and finding a job so that its justification is some economic truth or principle which is taken for granted (implicit)? Again, the way to proceed is to ask, 'What would *show* that,

everyone *could* solve their unemployment problem by great ingenuity in hunting for a job or by willingness to work for less?'

Presumably the way to show this is by means of sound economic argument, based on well-established economic truths. Since the argument we are considering does not do this – does not do what is *required*, in order to establish its conclusion, it is *not* a good argument and its conclusion (may be true but) does not follow from the reason given. Passage (8) is still an argument, the use of the word 'so' makes that clear, but its reason could be true and its conclusion false – the argument does not establish its conclusion. Interestingly, this argument is discussed by Paul Samuelson in his widely used and influential economics textbook *Economics: An Introductory Analysis* where he presents it in a list of classic economic *fallacies*.

These examples are enough to begin with. We now introduce a more general approach.

2 · A general method of argument analysis

In Chapter 1 we considered several *examples*: most of them were arguing a case and we used them to point up various lessons about reasoning. Having given the reader a taste of argument analysis we now introduce a *general* method for analysing and evaluating arguments. The method lay behind what was said in Chapter 1 but the reader who tried the exercises should now be ready for a general account rather than the piecemeal approach.

The method which we describe applies to reasoning, or argument, as it actually occurs in natural language – in our case, English. We begin by describing how to recognise contexts in which reasoning is taking place (i.e. we say what the 'linguistic clues' are). We then describe how to uncover and display the *structure* of a piece of reasoning (whether it is a 'chain' of reasons etc.). Finally we explain, as far as possible, how to decide whether the reasoning is correct or incorrect.

At this stage we do no more than *outline* the method. We do this so that its essential lines may be boldly drawn and clearly grasped. Too many qualifications at this point might obscure the method's basic simplicity: if it is basically correct the place to develop and refine it is where the problems arise – in applying it to particular examples – and this is what we shall do. In subsequent chapters the basic skeleton will be extended and 'fleshed out' as the need arises. We shall do this in the course of showing how to apply the method to a number of instructive examples.

Nearly all the arguments we study in this book are arguments which have *actually been used* by someone with a view to *convincing* others about some matter. They are *real* arguments – not the 'made-up' kind with which logicians usually deal. They originate from various sources ranging from classic texts to newspapers. And they come from various fields, although broadly speaking they are from the social sciences, some natural sciences and philosophy.

I The language of reasoning

Some clues

Of course we use language for many purposes besides reasoning. We use it to report events, to tell jokes, to extend invitations, to tell stories, to make

promises, to give orders, to ask questions, to issue instructions, to evoke emotions, to describe things, to entertain, and a thousand other things. (It is as well to note early on that in this book we are restricting our attention to a quite specific area of human activity, even if it is of very general importance.) Each of the activities mentioned above employs its own characteristic language – a language which helps us to grasp what is happening. For example ‘Have you heard the one about . . . ?’ is a common way to signal that a joke is coming (rather than a true report, etc.). ‘Would you like to come and . . . ?’ is a common way of issuing an invitation. ‘Don’t do that or else . . . !’ is a common way of issuing a threat, and so on. Of course these same phrases can be used for quite different purposes and knowing the context in which they are used is often essential to understanding their meaning. It would be a very complex matter to say how in general we identify jokes, or threats or whatever (cf. J. L. Austin, *How To Do Things With Words*) and it should come as no surprise that the language of reasoning is complex too, but there are several helpful things which can be said.

If we are to focus attention on reasoning we must first describe how to distinguish contexts in which reasoning is taking place. Remember, reasoning or arguing a case consists in giving grounds or *reasons* for *conclusions*, and the reasons are put forward in order to *support, justify, establish, prove* or *demonstrate* the conclusion. (The author is trying to *convince* the audience by means of reasoning.) In natural languages it is not always easy to tell when an argument is being presented (remember some of the examples in Chapter 1), but all arguments have a conclusion and in English a conclusion is often signalled by the presence of one of the following words or phrases, which we call ‘conclusion indicators’:

<i>Conclusion indicators</i>	
therefore	I conclude that
so which implies that
hence which allows us to
thus	infer that
consequently	it follows that
which proves that establishes the fact that
justifies the belief that demonstrates that

We do not suggest that *whenever* these phrases are used a conclusion follows, but that they commonly *indicate* the presence of a conclusion. They are linguistic clues to what is intended in the text. Sometimes of course they have a quite different usage, for example ‘I is thirty minutes since I started to read this book’, ‘I can only go so far’, ‘You ride a bicycle *thus*’. The conclusion indicators which we have listed, and similar ones, are only markers. They cannot be used mechanically to find conclusions: it usually requires a little judgement to decide whether they do in fact signal the presence of a conclusion. Of course, conclusions are sometimes presented with

no conclusion indicator at all, instead the context shows that a conclusion is being presented.

Every argument also includes grounds or *reasons* for its conclusion. A reason is usually presented as being *true* and as being *a reason for* some conclusion. (For the sake of simplicity we begin with examples in which the reasons are *presented as being true* and we restrict the term ‘reason’ to such cases. However, in Chapter 8 we shall extend the term to include reasons which are *not* presented as being true but which are ‘supposed for the sake of the argument’.) Words and phrases which are used in English to signal the presence of reasons – and which we shall call ‘reason indicators’ – include the following:

<i>Reason indicators</i>	
because	the reason being
for	firstly, secondly, (etc.)
since	may be inferred from the fact that
follows from the fact that	

Again, we are not saying that *whenever* these words or phrases are used a reason is present, but that they commonly *indicate* the presence of a reason. They serve as markers which enable us, with the aid of a little judgement, to locate reasons. Again it is true that reasons are often presented without reason indicators but that the context shows that a reason is being given.

It will be convenient to have a phrase to refer to both reason and conclusion indicators so we shall call them both ‘inference indicators’ or ‘argument indicators’.

Some complications

(i) The contexts in which we are interested are those in which an author or speaker presents some claim, the conclusion, as being supported or justified by other claims, the reasons. So whether certain claims are to be counted as conclusions or reasons depends solely on the author’s apparent intentions – as he or she has expressed them. It does *not* matter whether the claims are *true* or *false*, nor does it matter whether the reasons *succeed* in justifying the conclusion: all that matters at this stage – where we are trying to identify what the argument is – is whether the text presents some claims as *reasons for conclusions*.

(ii) Sometimes reasoning takes place without the use of inference indicators to signal the presence of reasons and conclusions. Sometimes it is difficult to decide whether reasoning is taking place in such cases. We shall explain shortly (pp. 22f.) how to make that decision. In general when trying to decide whether a passage contains reasoning one should adopt the Principle of Charity. This says that if interpreting as reasoning a passage

which is not *obviously* reasoning yields only bad arguments, assume it is *not* reasoning. (The rationale for this approach is that we are interested in finding out the truth about things rather than in scoring points off people.) Omitting inference indicators is sometimes a rhetorical device, used for purposes of emphasis, and is common with politicians and public speakers (cf. Weinberger's letter in Chapter 4).

(iii) There is an important complication which arises out of the different uses to which some inference indicators can be put. This is best explained by reference to the ambiguity of the word 'because' which sometimes signals the presence of a reason for a conclusion, but which sometimes signals the presence of a *causal claim* or, more generally, an *explanation* of some kind. Here are some examples.

- (1) John broke the window because he tripped.
- (2) John broke the window because he had forgotten his key.
- (3) John must have broken the window because he was the only person in the house.

Assuming the natural context in each case the question is how to understand what is being said. Clearly, in neither case (1) nor case (2) does the use of 'because' signal that a reason is being given for a conclusion. In (1) the whole statement is a *causal* one: what *caused* John to break the window was that he tripped. The whole statement could well be the *conclusion* of some other reasoning but in itself it expresses no argument at all. In (2) the statement *explains* John's *reason* for breaking the window — explains why he did it. Again the whole statement could be the conclusion of further reasoning but in itself it expresses no argument. In (3), on the other hand, the natural way to construe it requires that we treat 'because' as a reason indicator. (The 'must' is a further clue. See below.)

(iv) So-called 'modal' words and phrases like 'must', 'cannot', 'impossible', 'necessarily', and so on are sometimes used to signal reasoning. For example:

The engine won't fire. The carburettor must be blocked.

Assuming the obvious context, the word 'must' is being used by the speaker to signal the fact that he is drawing a conclusion: he could have said, 'Since the engine won't fire, I conclude that the carburettor is blocked' and this would have conveyed much the same message (though rather stuffily!). Here is another example:

The world is full of suffering. God cannot exist.

(v) The conclusion is sometimes omitted from an argument. Here is an example:

All boxers suffer brain damage and Smith had a long career as a boxer. (Need I say more?)

The context usually makes the intended conclusion clear. Similarly, reasons are sometimes omitted from arguments even though they are being assumed as part of the argument. We have discussed several such cases in Chapter 1 and we shall say more about them later.

It follows immediately from paragraphs (i) to (v) that the dividing line between argument and non-argument is not sharp. It is often absolutely clear that a passage expresses an argument. Similarly it is often quite certain that a passage does *not* contain reasoning. But equally, it is often quite unclear whether it does or not.

II The structure of reasoning

We have explained part of what is necessary in order to decide whether some piece of English contains reasoning, but in real life it is often surprisingly difficult to tell exactly what the argument is supposed to be, so we shall shortly describe a systematic and comprehensive method for extracting an argument from its text. Before we do this, however we shall find it helpful to introduce some conventions for representing arguments, some terminology, and some elementary ideas about the *structure* of an argument.

Some conventions and terminology

We begin by dealing with the *simplest* cases of reasoning, cases in which the reasons are presented as *being true* and as being reasons for some conclusion. (We shall extend the present treatment to deal with 'suppositional' arguments in Chapter 8.) The notation we are about to introduce is not *essential* to argument analysis: those who hate symbols may stick to words, like 'therefore' etc., but they do need to grasp the ideas behind the notation.

If some claim, R, is presented as being a reason for accepting some conclusion, C, we shall write it like this:

$$R \rightarrow C$$

where the arrow is to be read as 'therefore' or some idiomatically appropriate synonym.

If several reasons are given for some conclusion there are two possibilities: the reasons may be presented as *jointly* supporting the conclusion (taken *together* they support the conclusion but each *in isolation* does not) or they may be presented as *independently* justifying it (so that if you accept *one* of

the reasons the author expects you to accept the conclusion). An example of the latter case is the following:


Russia will not occupy Britain because she does not want to. Anyway, the Americans would not let her.

Another example is this:


Universities must expect further cuts because they have suffered less than other sectors of education, but even if this were not so, they should expect further cuts because they are not sufficiently vocationally oriented.

Examples where the reasons are presented as *jointly* supporting the conclusion were common in Chapter 1; (1), (2), (3) and (6) were just such examples.

Let us suppose that two reasons, R_1 and R_2 , are given for some conclusion C , then we shall represent the cases we have just described as follows. If R_1 and R_2 are *joint* reasons for C we write,




but if R_1 and R_2 are *independent* reasons for C we write,



(If it is not easy to judge which the author intends, choose whichever interpretation yields the *better* argument, i.e. whichever is the hardest to fault.)

This is all the notation we need for the moment. We shall extend it as we need to. Of course, complex arguments may combine the cases we have described in a variety of ways. In particular, the conclusion of one part of the argument may be used as a reason for some further conclusion (as in example (4), Chapter 1). We shall call such a conclusion an *intermediate conclusion*: it is presented both as a conclusion from prior reasons and as a reason for a subsequent conclusion. If a reason, R , is given for some conclusion, C , and the argument contains *no* intermediate conclusion between R and C , then we shall call R an *immediate reason* for C . Those reasons which are presented without themselves being supported by other reasons, we shall call the *basic reasons* or *premisses* of the argument. A conclusion which is not used in the argument to support any further conclusion will be called a *final conclusion* or *main conclusion*.

So, in summary, one might have an argument whose structure is represented by the following *argument diagram* (as we shall call it):



Basic reason₁ and basic reason₂ are not immediate reasons for the final conclusion, but intermediate conclusion₁ and basic reason₃ are.

Those who hate notation and diagrams can do everything using appropriate words and writing the reasoning in linear form, so the example diagrammed above might look as follows:

- (1) Basic reason₁.
- (2) Basic reason₂.
- Either (1) or (2) is sufficient to justify (3), therefore
- (3) Intermediate conclusion₁.
- (3) is true and
- (4) Basic reason₃.
- Therefore
- (5) Final conclusion.

Given a piece of reasoning in ordinary English, it is helpful either to draw its *argument diagram* or to write it in equivalent linear form. Doing so forces us to clarify exactly what the argument is – and that is its main purpose. It is also true however that the structure which is revealed by this process may be important in deciding whether the argument is sound. We shall return to this point later but we have now progressed far enough to outline a method for extracting an argument from its context, so we do this next.

The method of extracting arguments outlined

It is easy to underestimate how difficult it can be to extract an author's intended argument from a written, natural language text, but given a piece of ordinary English the following method will help to determine its conclusion(s), its reason(s) and the structure of its argument(s) if these are not already clear.

- (1) Read through the text to get its sense, circling – [thus] – all the inference indicators as you go.
- (2) Underline – thus – any clearly indicated conclusions, and bracket – (thus) – any clearly indicated reasons. (It helps at this stage if one tries to summarise the argument.)
- (3) Identify what you take to be the main conclusion and mark it C. (There may be more than one.)

- (4) Starting with C, ask 'What immediate reasons are presented in the text for accepting C?' or 'Why (in the text) am I asked to believe C?' Use inference indicators to help answer the question. If the question is hard to answer because the author's intentions are not transparent (i.e. they are neither explicitly shown by argument indicators nor obvious from the context), then ask the Assertibility Question.

(AQ) What argument or evidence would justify me in asserting the conclusion C? (What would I have to know or believe to be justified in accepting C?)

Having done this look to see if the author asserts or clearly assumes these same claims (reasons). If he does it is reasonable (and accords with the Principle of Charity) to construe him as having intended the same argument. If he doesn't you have no rational way of reconstructing his argument (on the basis of the text alone).

- (5) For each reason, R, already identified, repeat the process described in step (4) above. Do this until you are left with only basic reasons and then display the argument(s) in a clear way (say, by means of a diagram or in linear form).

This is the basic outline of the method, but several further points need to be made. If it is to be properly understood. We present some general ones first and reserve some specific ones to the next section.

- (a) Notice that the issue is 'What does the text/author present as a reason, conclusion, etc.?' not 'What is a good reason, etc.?' But notice also that in order to find the answer to the first question we may have to ask ourselves the second one. To put the point another way,
- (b) inference indicators may make an author's intentions completely clear (quite certain); context may do the same; but if this is not the case the only way you can divine the author's intentions (given only the text) is to construct the *best* argument you can and ask whether the author could be construed as presenting it. It follows that in such cases,
- (c) this is not a mechanical method which yields an argument automatically; it requires judgement and imagination.
- (d) Furthermore, the extent to which you can grasp the author's intended meaning will depend on your understanding of the language and your knowledge of the subject and so will be a matter of degree.
- (e) The philosophical justification for the use of the Assertibility Question is based on the assumption that

- * If you understand a proposition you must be able to give at least some account of how you could decide whether it was true or false, what argument or evidence would show it to be true or false (otherwise you don't understand it at all).
We shall return to this principle often in what follows, but we leave it unsupported for the moment.

The structure of reasons and conclusions

Just as arguments are logically complex, so are their reasons and conclusions: they too exhibit logical structure. We shall say very little here about these 'internal' structures (see the Appendix for a fuller account) but some things must be said in order to avoid confusion in applying the method just outlined. In short there are cases where the internal structure of a sentence may obscure what is to *count* as a reason or a conclusion.

In order to discuss these cases we shall find it convenient to have just one word to describe the case when an author claims that something is true (presents it as being true); we shall say that such a claim is *asserted* and we shall call it an *assertion*. (For a fuller explanation, see Appendix, p. 173.)

Suppose that police have evidence leading them to the conclusion that,

Either Jones killed Brown, or Smith did (J or S)

For present purposes, the important thing to notice is that the police are *not* asserting that Jones did it, *nor* are they asserting that Smith did it: they are asserting the whole 'disjunction' (as logicians call it) 'J or S'. So in identifying reasons and conclusions, disjunctions must not be broken up into their parts. (Of course, if Smith produces a cast-iron alibi, this, together with the police conclusion 'J or S', yields an argument with the *asserted* conclusion, 'Jones killed Brown' – but that is a further stage in the argument.)

Disjunctions don't often create problems in argument analysis, but hypotheticals do. Remember that a hypothetical is a sentence of the form 'if... then...'. Clearly, when someone says,

+ If [the money supply is increasing in Britain], then {the rate of inflation in Britain will increase}

they are *neither* asserting that the money supply is increasing in Britain *nor* asserting that the rate of inflation in Britain will increase. They are asserting the whole hypothetical, 'If [A] then [B]'. Such hypotheticals are very common and very important in reasoning. Partly because of this the 'parts' of a hypothetical have special names, the [A] part, the part governed by 'if', is called the 'antecedent' and the [B] part the 'consequent'. (In the hypothetical 'B, if A' the B part is still the 'consequent' etc.) There are two points we need to make about hypotheticals here.

Firstly, a hypothetical may occur as either a reason or a conclusion in the course of a piece of reasoning (e.g. ⁺ above could). In that case the hypothetical *should not be broken up* into antecedent and consequent. What is functioning as a reason, or conclusion – what is being asserted – is the whole hypothetical. Here is an example where the three reasons and the conclusion are all hypotheticals.

If every event has a cause then all my actions are caused. If all my actions are caused, I am not free to do what I choose. In that case I am not responsible for my actions. So if every event has a cause, I am not responsible for my actions.

Secondly, there are numerous phrases which signal the presence of a hypothetical statement: we mention some of them now (and call them 'hypothetical indicators').

Hypothetical indicators

- if ... then ...
- suppose ... then ...
- unless ... then ...
- ... provided that ...
- ... on the condition that ...

As with inference indicators, these phrases are markers which have to be used with understanding.

Much more could be said about the internal logical structure of sentences but we have said all we need for the present. Further explanation is provided in the Appendix and there are many excellent logic texts which develop the subject in various ways, but more detail is unnecessary for our purposes at this stage.

This completes our outline of the method for extracting an argument from its context; we now move on to the method for evaluating it.

III Tests for a good argument

Once it is clear what argument we are considering then we are in a position to test whether it establishes its conclusion. Remember that we are still restricting our attention to those arguments in which the reasons are presented as being true.

In real arguments the first thing people normally challenge is the truth of the premises. If the premises of an argument are not true (or at least *one* must be true in the case of independent reasons) then they cannot establish their conclusion, so the argument loses much of its interest. (Although in

theoretical arguments – or in contexts where you don't know if the premises are true – it may still be interesting to ask, 'If the premises *were* true would they establish the conclusion?') So the first condition an argument must satisfy if it is to establish its conclusion is,

I All its premises must be true – except that where independent reasons are given for some conclusion at least *one* must be true.

Now let us suppose that the argument we are considering has true premises, so that it satisfies condition I. At this point in real arguments, people who think the argument fails to establish its conclusion will say things like, 'the conclusion is *not justified*' or 'the argument is *not sound*' or 'the conclusion *does not follow* from the premises'. It is easy to think of an argument which has true premises but whose conclusion does not follow, for example,

- (1) All women are mortal
- and (2) The American President is mortal
- therefore
- (3) The American President is a woman.



(Ask yourself why the conclusion does *not* follow in this example. Does the conclusion follow in 'All men are mortal and the American President is mortal so the American President is a man?')

So the second condition an argument must satisfy if it is to establish its conclusion is,

II The conclusion must *follow* from the premises

and we must now explain how one decides whether this is the case.

The intuitive idea is this: a conclusion *follows from* its premises if and only if the truth of the premises *guarantees* the truth of the conclusion, so the test to apply is,

Could the premises be true and the conclusion false?

If the answer to the question is 'Yes' the conclusion does *not* follow from its premises. If the answer is 'No' then the truth of the premises guarantees the truth of the conclusion, the conclusion *follows from* its premises – and if you accept the premises you must accept the conclusion.

An illustrative example

To illustrate how the test works consider an example. Let A, B and C be politicians or policies of your choice and ask whether the conclusion follows in this argument:

enter into the realm of philosophy, in particular the 'theory of knowledge', though you need no expertise in formal philosophy. We shall show how the method works out in a number of instructive examples in the rest of the book. Much of our discussion will be about choosing 'appropriate standards': these are not established, objective facts, nor are they arbitrary, they require a good deal of argument. The concluding chapter will then attempt to provide a philosophical justification for the method outlined here and applied in the intervening chapters.

3 • A first example – from Thomas Malthus

If the present world population doubles itself every twenty-five years, in 150 years' time there will be standing room only since the number of people will be greater than the number of square metres on the land surface of the Earth.

In this chapter we show how to analyse and evaluate a very famous argument due to Thomas Malthus (1766–1834) and we apply and develop the method of Chapter 2 in the process. Malthus's father was a friend of David Hume and of Jean-Jacques Rousseau, both of whom visited his house together when Thomas was only three weeks old. It was under the influence of Rousseau's *Emile* that his father had Thomas privately educated until he became an undergraduate at Jesus College, Cambridge, at the age of eighteen, in 1784. He graduated well in mathematics in 1788, and he took Holy Orders in the same year. His *Essay on the Principle of Population as it affects the Future Improvement of Society with Remarks on the Speculations of Mr. Godwin, M. Condorcet and other Writers* was first published in 1798. There was much discussion at that time – in the wake of the French Revolution – about the possibility of establishing a society based on social and economic equality. Malthus's *Essay* originated as a polemic against such utopian speculations. His argument was not new.

The most important argument that I shall adduce is certainly not new. The principles on which it depends have been explained in part by Hume, and more at large by Dr. Adam Smith.

We now present Malthus's basic argument. Our extract consists of most of Chapter 1 of the *Essay* and is the part which Malthus explicitly described in his text as 'an outline of the principal argument of the essay'.

I Thomas Malthus, *An Essay on the Principle of Population*, Chapter 1

(Successive paragraphs are labelled to enable easy reference to them later.)

- (a) In entering upon the argument I must premise that I put out of the question, at present, all mere conjectures, that is, all suppositions,