1 Basic Tools for Argument

1.1 Arguments, premises, and conclusions

Philosophy is for nit-pickers. That's not to say it is a trivial pursuit. Far from it. Philosophy addresses some of the most important questions human beings ask themselves. The reason philosophers are nit-pickers is that they

The Philosopher's Toolkit: A Compendium of Philosophical Concepts and Methods, Third Edition. Peter S. Fosl and Julian Baggini.

are commonly concerned with the ways in which the claims and beliefs people hold about the world either are or are not rationally supported, usually by rational argument. Because their concern is serious, it is important for philosophers to demand attention to detail. People reason in a variety of ways using a number of techniques, some legitimate and some not. Often one can discern the difference between good and bad reasoning only if one scrutinises the content and structure of arguments with supreme and uncompromising diligence.

Argument and inference

What, then, is an 'argument' proper? For many people, an argument is a contest or conflict between two or more people who disagree about something. An argument in this sense might involve shouting, name-calling, and even a bit of shoving. It might also – but need not – include reasoning.

Philosophers, in contrast, use the term 'argument' in a very precise and narrow sense. For them, an argument is the most basic complete unit of reasoning – an atom of reasoning. An 'argument' understood this way is an *inference* from one or more starting points (truth claims called a 'premise' or 'premises') to an end point (a truth claim called a 'conclusion'). All arguments require an inferential movement of this sort. For this reason, arguments are called *discursive*.

Argument vs explanation

'Arguments' are to be distinguished from 'explanations'. A general rule to keep in mind is that arguments attempt to demonstrate *that* something is true, while explanations attempt to show *how* something is true. For example, consider encountering an apparently dead woman. An explanation of the woman's death would undertake to show *how* it happened. ('The existence of water in her lungs explains the death of this woman.') An argument would undertake to demonstrate *that* the person is in fact dead ('Since her heart has stopped beating and there are no other vital signs, we can conclude that she is in fact dead.') or that one explanation is better than another ('The absence of bleeding from the laceration on her head combined with water in the lungs indicates that this woman died from drowning and not from bleeding.')

The place of reason in philosophy

It's not universally realised that reasoning comprises a great deal of what philosophy is about. Many people have the idea that philosophy is essentially about ideas or theories about the nature of the world and our place in it that amount just to *opinions*. Philosophers do indeed advance such ideas and theories, but in most cases their power, their scope, and the characteristics that distinguish them from mere opinion stem from their having been derived through rational argument from acceptable premises. Of course, many other regions of human life also commonly involve reasoning, and it may sometimes be impossible to draw clean lines demarcating philosophy from them. (In fact, whether or not it is possible to demarcate philosophy from non-philosophy is itself a matter of heated philosophical debate!)

The natural and social sciences are, for example, fields of rational inquiry that often bump up against the borders of philosophy (especially in inquiries into the mind and brain, theoretical physics, and anthropology). But theories composing these sciences are generally determined through certain formal procedures of experimentation and reflection to which philosophy has little to add. Religious thinking sometimes also enlists rationality and shares an often-disputed border with philosophy. But while religious thought is intrinsically related to the divine, sacred, or transcendent – perhaps through some kind of revelation, article of faith, or ritualistic practice – philosophy, by contrast, in general is not.

Of course, the work of certain prominent figures in the Western philosophical tradition presents decidedly non-rational and even anti-rational dimensions (for example, that of Heraclitus, Kierkegaard, Nietzsche, Heidegger, and Derrida). We will examine the non-argumentative philosophical methods of these authors in what follows of this book. Furthermore, many include the work of Asian (Confucian, Taoist, Shinto), African, Aboriginal, and Native American thinkers under the rubric of philosophy, even though they seem to make little use of argument and have generally not identified their work as philosophical.

But, perhaps despite the intentions of its authors, even the work of non-standard thinkers involves rationally justified claims and subtle forms of argumentation too often missed. And in many cases, reasoning remains on the scene at least as a force with which thinkers must reckon.

Philosophy, then, is not the only field of thought for which rationality is important. And not all that goes by the name of philosophy is argumentative. But it is certainly safe to say that one cannot even begin to master the expanse of philosophical thought without learning how to use the tools of reason. There is, therefore, no better place to begin stocking our philosophical toolkit than with rationality's most basic components, the subatomic particles of reasoning – 'premises' and 'conclusions'.

Premises and conclusions

For most of us, the idea of a 'conclusion' is as straightforward as a philosophical concept gets. A conclusion is just that with which an argument concludes, the product and result of an inference or a chain of inferences, that which the reasoning claims to justify and support. What about 'premises', though? Premises are defined in relation to the conclusion. They are, of course, what do the justifying. There is, however, a distinctive and a bit less obvious property that all premises and conclusions must possess.

In order for a sentence to serve either as a premise or as a conclusion, it must exhibit this essential property: it must make a claim that is either true or false. A sentence that does that is in logical terms called a *statement* or *proposition*.

Sentences do many things in our languages, and not all of them possess that property and thence not all of them are statements. Sentences that issue commands, for example ('Forward march, soldier!'), or ask questions ('Is this the road to Edinburgh?'), or register exclamations ('Wow!'), are neither true nor false. Hence, it's not possible for sentences of those kinds to serve as premises or as conclusions.

This much is pretty easy, but things can get sticky in a number of ways. One of the most vexing issues concerning arguments is the problem of implicit claims. That is, in many arguments, key premises or even the conclusion remain unstated, implied or masked inside other sentences. Take, for example, the following argument: 'Socrates is a man, so Socrates is mortal.' What's left implicit is the claim that 'all men are mortal'. Arguments with unstated premises like this are often called *enthymemes* or *enthymemetic*.

It's also the case that sometimes arguments nest inside one another so that in the course of advancing one, main conclusion several ancillary conclusions are proven along the way. Untangling arguments nested in others can get complicated, especially as those nests can pile on top of one another and interconnect. It often takes a patient, analytical mind to sort it all out (just the sort of mind you'll encounter among philosophers).

In working out precisely what the premises are in a given argument, then, ask yourself first what the principal claim is that the argument is trying to demonstrate. Then ask yourself what other claims the argument relies upon (implicitly or explicitly) in order to advance that demonstration. Sometimes certain words and phrases will explicitly indicate premises and conclusions. Phrases like 'therefore', 'in conclusion', 'it follows that', 'we must conclude that', and 'from this we can see that' often indicate conclusions. ('The DNA, the fingerprints, and the eyewitness accounts all point to Smithers. It follows that she must be the killer.') Words like 'because' and 'since', and phrases like 'for this reason' and 'on the basis of this', on the other hand, often indicate premises. (For example, 'Since the DNA, the fingerprints, and the eyewitness accounts all implicate Smithers, she must be the killer.')

Premises of an argument, then, compose the set of claims from which the conclusion is drawn. In other sections, the question of precisely how we can justify the move from premises to conclusion will be addressed in more in more detail (see 1.4 and 4.7). But before we get that far, we must first ask, 'What justifies a reasoner in entering a premise in the first place?'

[redacted]

[redacted]

(6)

SEE ALSO

1.10 Definitions
3.7 Circularity
7.1 Basic beliefs
7.9 Self-evident truths

READING

★ Nigel Warburton (2000). *Thinking From A to Z*, 2nd edn
John Shand (2000). *Arguing Well*
★ Graham Priest (2001). *Logic: A Very Short Introduction*
Peter Klein (2008). 'Contemporary responses to Agrippa's trilemma' in *The Oxford Handbook of Skepticism* (ed. John Greco)

1.2 Deduction

The murder was clearly premeditated, Watson. The only person who knew where Dr Fishcake would be that night was his colleague, Dr Salmon. Therefore, the killer must be …

Deduction is the form of reasoning that is often emulated in the formulaic drawing-room denouements of classic detective fiction. It is the most rigorous form of argumentation there is, since in deduction the move from premises to conclusions is such that if the premises are true, then the conclusion *must* (*necessarily*) also be true. For example, take the following argument:

1. Elvis Presley lives in a secret location in Idaho.
2. All people who live in secret locations in Idaho are miserable.
3. Therefore, Elvis Presley is miserable.

If we look at our definition of a deduction, we can see how this argument fits the bill. If the two premises are true, then the conclusion must also definitely be true. How could it not be true that Elvis is miserable, if it is indeed true that all people who live in secret locations in Idaho are miserable, and Elvis is one of those people?

You might well be thinking there's something fishy about this, since you may believe that Elvis is not miserable for the simple reason that he no longer exists. So, all this talk of the conclusion having to be true might strike you as odd. If this is so, you haven't taken on board the key word at the start of this sentence, which does such vital work in the definition of deduction. The conclusion must be true *if* the premises are true. This is a big 'if'. In our example, the conclusion is, we confidently believe, not true and for very good reasons. But that doesn't alter the fact that this is a deductive argument, since if it turned out that Elvis does live in a secret location in Idaho and that all people who lived in secret locations in Idaho are miserable, it would necessarily follow that Elvis is miserable.

The question of what makes a good deductive argument is addressed in more detail in the section on validity and soundness (1.4). But in a sense, everything that you need to know about a deductive argument is contained within the definition just given: a (successful) deductive argument is one where, if the premises are true, then the conclusion is definitely true.

Before we leave this topic, however, we should return to the investigations pursued by our detective. Reading his deliberations, one could easily insert the vital, missing words. The killer must surely be Dr Salmon. But is this the conclusion of a successful deductive argument? The fact is that we can't answer this question unless we know a little more about the exact meaning of the premises.

First, what does it mean to say the murder was 'premeditated'? It could mean lots of things. It could mean that it was planned right down to the last detail, or it could mean simply that the murderer had worked out what she would do in advance. If it is the latter, then it is possible that the murderer did not know where Dr Fishcake would be that night, but, coming across him by chance, put into action her premeditated plan to kill him. So, it could be the case (1) that both premises are true (the murder was premeditated, and Dr Salmon was the only person who knew where Dr Fishcake would be that night) but (2) that the conclusion is false (Dr Salmon is, in fact, not the murderer). Therefore, the detective has not formed a successful deductive argument.

What this example shows is that, although the definition of a deductive argument is simple enough, spotting and constructing successful deductive arguments is much trickier. To judge whether or not the conclusion really *must* follow from the premises, you have to be sensitive to ambiguity in the premises as well as to the danger of accepting too easily a conclusion that seems to be supported by the premises but does not in fact follow from them. Deduction is not about jumping to conclusions, but crawling (though not slouching) slowly towards them.

SEE ALSO

1.1 Arguments, premises, and conclusions
1.3 Induction
1.4 Validity and soundness

READING

★ Alfred Tarski (1936/95). *Introduction to Logic and to the Methodology of Deductive Sciences*
★ Fred R. Berger (1977). *Studying Deductive Logic*
★ A.C. Grayling (2001). *An Introduction to Philosophical Logic*
Warren Goldfarb (2003). *Deductive Logic*
★ Maria Konnikova (2013). *Mastermind: How to Think Like Sherlock Holmes*

1.3 Induction

I (Julian Baggini) have a confession to make. Once, while on holiday in Rome, I visited the famous street market, Porta Portese. I came across a man who was taking bets on which of the three cups he had shuffled around was covering a die. I will spare you the details and any attempts to justify my actions on the grounds of mitigating circumstances. Suffice it to say, I took a bet and lost. Having been budgeted so carefully, the cash for that night's pizza went up in smoke.

My foolishness in this instance is all too evident. But is it right to say my decision to gamble was 'illogical'? Answering this question requires wrangling with a dimension of logic philosophers call 'induction'. Unlike deductive inferences, induction involves an inference where the conclusion follows from the premises not with necessity or definitely but only with *probability* (though even this formulation is problematic, as we'll see).

Defining induction

Perhaps most familiar to people is a kind of induction that involves reasoning from a limited number of observations to wider generalisations of some probability. Reasoning this way is commonly called *inductive generalisation*. It's a kind of inference that usually involves reasoning from past regularities to future regularities. One classic example is the sunrise. The sun has risen regularly each day, so far as human experience can recall, so people reason that it will probably rise tomorrow. This sort of inference is often taken to typify induction. In the case of my Roman holiday, I might have reasoned that the past experiences of people with average cognitive abilities like mine show that the probabilities of winning against the man with the cups is rather small.

But beware: *induction is not essentially defined as reasoning from the specific to the general.* An inductive inference need not be past–future directed. And it can involve reasoning from the general to the specific, the specific to the specific, or the general to the general.

I could, for example, reason from the *more general*, past-oriented claim that no trained athlete on record has been able to run 100 metres in under 9 seconds, to the *more specific* past-oriented conclusion that my friend had probably not achieved this feat when he was at university, as he claims. Reasoning through *analogies* (see 2.4) as well as *typical examples* and *rules of thumb* are also species of induction, even though none of them involves moving from the specific to the general. The important property of inductive inferences is that they determine conclusions only with probability, not how they relate specific and general claims.

The problem of induction

Although there are lots of kinds of induction besides inductive generalisations, that species of induction is, when it comes to actual practices of reasoning, often where the action is. Reasoning in experimental science, for example, commonly depends on inductive generalisations in so far as scientists formulate and confirm universal natural laws (e.g. Boyle's ideal gas law) only with a degree of probability based upon a relatively small number of observations. Francis Bacon (1561–1626) argued persuasively for just this conception of induction.

The tricky thing to keep in mind about inductive generalisations, however, is that they involve reasoning from a 'some' in a way that in deduction would require an 'all' (where 'some' means at least one but perhaps not all of some set of relevant individuals). Using a 'some' in this way makes inductive generalisation fundamentally different from deductive argument (for which such a move would be illegitimate). It also opens up a rather enormous can of conceptual worms. Philosophers know this conundrum as the *problem of induction*. Here's what we mean. Take the following example:

1. *Almost all* elephants like chocolate.
2. This is an elephant.
3. Therefore, this elephant likes chocolate.

This is *not* a well-formed deductive argument, since the premises could possibly be true and the conclusion still be false. Properly understood, however, it may be a strong inductive argument – if the conclusion is taken to be probable, rather than certain.

On the other hand, consider this rather similar argument:

1. *All* elephants like chocolate.
2. This is an elephant.
3. Therefore, this elephant likes chocolate.

Though similar in certain ways, this one is, in fact, a well-formed deductive argument, not an inductive argument at all. One way to think of the problem of induction, therefore, is as the problem of how an argument can be good reasoning as induction but be poor reasoning as a deduction. Before addressing this problem directly, we must take care not to be misled by the similarities between the two forms.

A misleading similarity

Because of the general similarity one sees between these two arguments, inductive arguments can sometimes be confused with deductive arguments. That is, although they may actually look like deductive arguments, some arguments are actually inductive. For example, an argument that the

D.C. Stove (1986/2001). *The Rationality of Induction*
★ Colin Howson (2003). *Hume's Problem: Induction and the Justification of Belief*

1.4 Validity and soundness

In his book, *The Unnatural Nature of Science*, the eminent British biologist Lewis Wolpert (b. 1929) argued that the one thing that unites almost all of the sciences is that they often fly in the face of common sense. Philosophy, however, may exceed even the (other?) sciences on this point. Its theories, conclusions, and terms can at times be extraordinarily counterintuitive and contrary to ordinary ways of thinking, doing and speaking.

Take, for example, the word 'valid'. In everyday speech, people talk about someone 'making a valid point' or 'having a valid opinion'. In philosophical speech, however, the word 'valid' is reserved exclusively for arguments. More surprisingly, a valid argument can look like this:

1. All blocks of cheese are more intelligent than any philosophy student.
2. Meg the cat is a block of cheese.
3. Therefore, Meg the cat is more intelligent than any philosophy student.

All utter nonsense, you may think, but from a strictly logical point of view this is a perfect example of a valid argument. How can that be so?

Defining validity

Validity is a property of well-formed deductive arguments, which, to recap, are defined as arguments where the conclusion in some sense (actually, hypothetically, etc.) follows from the premises *necessarily* (see 1.2). Calling a deductive argument 'valid' affirms that the conclusion actually does follow from the premises in that way. Arguments that are presented as or taken to be successful deductive arguments, but where the conclusion does not in fact definitely follow from the premises, are called 'invalid' deductive arguments.

The tricky thing, in any case, is that an argument may possess the property of validity even if its premises or its conclusion are *not* in fact *true*.

Validity, as it turns out, is essentially a property of an argument's *structure* or *form*; and so, the *content* and *truth value* of the statements composing the argument are irrelevant. Let's unpack this.

Consider structure first. The argument featuring cats and cheese given above is an instance of a more general argumentative structure, of the form:

1. All Xs are Ys.
2. Z is an X.
3. Therefore, Z is a Y.

In our example, 'block of cheese' is substituted for X, 'things that are more intelligent than all philosophy students' for Y, and 'Meg' for Z. That makes our example just one particular instance of the more general argumentative form expressed with the variables X, Y, and Z.

What you should notice is that you don't need to attach any particular meaning to the variables for this particular form to be a valid one. No matter with what we replace the variables, it will always be the case that *if* the premises are true (even though in fact they might not be), the conclusion *must* also be true. If there's *any* conceivable way possible for the premises of an argument to be true but its conclusion simultaneously be false, any coherent way at all, then it's an invalid argument.

This boils down to the notion of validity as content-blind or *topic-neutral*. It really doesn't matter what the content of the propositions in the argument is – validity is determined by the argument having a solid, deductive structure. Our block-of-cheese example is then a valid argument, because *if* its ridiculous premises were true, the ridiculous conclusion would also have to be true. The fact that the premises are ridiculous is neither here nor there when it comes to assessing the argument's validity.

The truth machine

Another way of understanding how arguments work as to think of them along the model of sausage machines. You put ingredients (premises) in, and then you get something (conclusions) out. Deductive arguments may be thought of as the best kind of sausage machine because they *guarantee* their output in the sense that when you put in entirely good ingredients (all true premises), you get out a fine-quality product (true conclusions). Of course, if you don't start with good ingredients, deductive arguments don't guarantee a good end product.

Invalid arguments are not generally desirable machines to employ. They provide no guarantee whatsoever for the quality of the end product. You might put in good ingredients (true premises) and sometimes get a high-quality result (a true conclusion). Other times good ingredients might yield a frustratingly poor result (a false conclusion).

Stranger still (and very different from sausage machines), with invalid deductive arguments you might sometimes put in poor ingredients (one or more false premises) but actually end up with a good result (a true conclusion). Of course, in other cases with invalid machines you put in poor ingredients and end up with rubbish. The thing about invalid machines is that you don't know what you'll get out. With valid machines, when you put in good ingredients (though *only* when you put in good ingredients), you have assurance. In sum:

Invalid argument
Put in false premise(s) → get out either a true or false conclusion
Put in true premise(s) → get out either a true or false conclusion

Valid argument
Put in false premise(s) → get out either a true or false conclusion
Put in true premise(s) → get out always and only a true conclusion

Soundness

To say an argument is valid, then, is not to say that its conclusion must be accepted as true. The conclusion is definitely established as true *only if* both of two conditions are met: (1) the argument is valid *and* (2) the premises are true. This combination of valid argument plus true premises (and therefore a true conclusion) is called approvingly a *sound* argument. Calling it sound is the highest endorsement one can give an argument. If you accept an argument as sound, you are really saying that one must accept its conclusion. The idea of soundness can even itself be formulated as an especially instructive valid, deductive argument:

1. If the premises of the argument are true, then the conclusion must also be true (i.e. the argument is valid).
2. The premises of the argument are true.
3. Therefore, the conclusion of the argument must also be true.

For a deductive argument to pass muster, it must be valid. But being valid is by itself not sufficient to make it a sound argument. A sound argument must not only be valid; it must have true premises, as well. It is, strictly speaking, only sound arguments whose conclusions we *must* accept.

Importance of validity

This may lead you to wonder why, then, the concept of validity has any importance. After all, valid arguments can be absurd in their content and false in their conclusions – as in our cheese and cats example. Surely it is soundness that matters?

Okay, but keep in mind that validity is a required component of soundness, so there can be no sound arguments without valid ones. Working out whether or not the claims you make in your premises are true, while important, is also not enough to ensure that you draw true conclusions. People make this mistake all the time. They forget that one can begin with a set of entirely true beliefs but reason so poorly as to end up with entirely false conclusions. It can be crucial to remember that starting with truth doesn't guarantee ending up with it.

Furthermore, for the sake of launching criticisms, it is important to grasp that understanding validity gives you an additional tool for evaluating another's position. In criticising a specimen of reasoning, you can either:

1. attack the truth of the premises from which he or she reasons,
2. or show that his or her argument is invalid, regardless of whether or not the premises deployed are true.

Validity is, simply put, a crucial ingredient in arguing, criticising, and thinking well, even if not the only ingredient. It's an utterly indispensable philosophical tool. Master it.

SEE ALSO

1.1 Arguments, premises, and conclusions
1.2 Deduction
1.5 Invalidity

READING

Aristotle (384–322 BCE). *Prior Analytics*
Fred R. Berger (1977). *Studying Deductive Logic*
S.K. Langer (2011). 'Truth and validity'. In: *Introduction to Symbolic Logic*, 3rd edn, Ch. 1, pp. 182–90
★ Jc Beall and Shay Allen Logan (2017). *Logic: The Basics*, 2nd edn

1.5 Invalidity

Given the definition of a valid argument, it may seem obvious what an invalid one looks like. Certainly, it's simple enough to define an invalid argument: it is an argument where the truth of the premises does not guarantee the truth of the conclusion. To put it another way, if the premises of an invalid argument are true, the conclusion may still be false. Invalid arguments are unsuccessful deductions and therefore, in a sense, are not truly deductions at all.

To be armed with an adequate definition of invalidity, however, may not be enough to enable you to make use of this tool. The man who went looking for a horse equipped only with the definition 'solid-hoofed, herbivorous, domesticated mammal used for draught work and riding' (*Collins English Dictionary*) discovered as much, to his cost. In addition to the definition, you need to understand the definition's full import. Consider this argument:

1. Vegetarians do not eat pork sausages.
2. Gandhi did not eat pork sausages.
3. Therefore, Gandhi was a vegetarian.

If you're thinking carefully, you'll have probably noticed that this is an invalid argument. But it wouldn't be surprising if you and a fair number of readers required a double take to see that it is in fact invalid. Now, this is a clear case, and if a capable intellect can easily miss a clear case of invalidity in the midst of an article devoted to a careful explanation of the concept, imagine how easy it is not to spot invalid arguments more generally.

One reason why many will fail to notice that this argument is invalid is because all three propositions are true. If nothing false is asserted in the premises of an argument and the conclusion is true, it's easy to think that the argument is therefore valid (and sound). But remember that an argument is valid *only if* the truth of the premises *guarantees* the truth of the conclusion in the sense that because of the argument's structure the conclusion is never false when the premises are true. In this example, this isn't so. After all, a person may not eat pork sausages yet not be a vegetarian. He or she may, for example, be an otherwise carnivorous Muslim or Jew. He or she simply may not like pork sausages but frequently enjoy turkey or beef.

So, the fact that Gandhi did not eat pork sausages does *not*, in conjunction with the first premise, guarantee that he was a vegetarian. It just so happens that he was. But, of course, since an argument can only be sound if it's valid, the fact that all three of the propositions it asserts are true does *not* make it a sound argument.

Remember that validity is a property of an argument's structure of form. In this case, the form is:

1. All Xs are Ys.
2. Z is a Y.
3. Therefore, Z is an X.

Here X is substituted for 'vegetarian', Y for 'person who does not eat pork sausages', and Z for 'Gandhi'. We can see why this structure is invalid by replacing these variables with other terms that produce true premises but a clearly false conclusion. (Replacing terms creates what logicians call a new '*substitution instance*' of the argument form.) If we substitute 'cat' for X, 'meat eater' for Y, and 'the president of the United States' for Z, we get:

1. All cats are meat eaters.
2. The president of the United States is a meat eater.
3. Therefore, the president of the United States is a cat.

The premises are true, but the conclusion clearly false. This cannot therefore be a valid argument structure. (Showing that an argument form is invalid by making substitutions that result in true premises but a false conclusion is called *showing invalidity by 'counterexample'*. It's a powerful skill well worth cultivating. See 1.7 and 3.12.)

It should be clear now that, as with validity, invalidity is not determined by the truth or falsehood of the premises but by the logical relations among them. This reflects a wider, and very important, feature of philosophy.

Philosophy is not just about saying things that are true or wise; it's about making true claims that are grounded in solid arguments. You may have a particular viewpoint on a philosophical issue, and it may just turn out by sheer luck that you're right. But, in many cases, unless you can demonstrate that you're right through good arguments, your viewpoint is not going to carry any weight in philosophy. Philosophers are not just concerned with the truth, but with what makes it the truth and how we can show that it's the truth.

SEE ALSO

1.2 Deduction
1.4 Validity and soundness
1.7 Fallacies

READING

★ Irving M. Copi (2010). *Introduction to Logic*, 14th edn
★ Harry Gensler (2016). *Introduction to Logic*, 3rd edn
★ Patrick J. Hurley and Lori Watson (2017). *A Concise Introduction to Logic*, 13th edn

1.6 Consistency

Ralph Waldo Emerson (1803–82) may have written in his well-known 1841 essay, 'Self-reliance,' that 'a foolish consistency is the hobgoblin of little minds,' but of all the philosophical crimes there are, the one with which you really don't want to get charged is inconsistency. For most purposes it's not too much to say that consistency is the cornerstone of rationality. To do philosophy well, therefore, it's crucial to master the idea and the practice of consistency.

Consistency is a property characterising two or more statements. If you hold two or more inconsistent beliefs, then, at root, this means you face a logically insurmountable problem with their truths. More precisely, the statements of your beliefs will be found to be somehow either to 'contradict'

[redacted]

SEE ALSO

1.12 Tautologies, self-contradictions, and the law of non-contradiction
2.1 Abduction
3.10 Contradiction/contrariety
7.2 Gödel and incompleteness
7.6 Paradoxes

READING

David Hilbert (1899). *Grundlagen der Geometrie*
★ P.F. Strawson (1952/2011). *Introduction to Logical Theory*
★ Fred R. Berger (1977). *Studying Deductive Logic*
★ Julian Baggini and J. Stangroom (2006). *Do You Think What You Think You Think?*
★ Aladdin M. Yaqub (2013). *Introduction to Logical Theory*

1.7 Fallacies

The notion of 'fallacy' will be an important instrument to draw from your toolkit, for philosophy often depends upon identifying poor reasoning, and a fallacy is nothing other than an instance of poor reasoning – a faulty inference. Since every invalid argument involves a faulty inference, a great deal of what one needs to know about fallacies has already been covered in the entry on invalidity (1.5). But while all invalid arguments are fallacious, not all fallacies involve invalid arguments. Invalid arguments are faulty because of flaws in their form or structure. Sometimes, however, reasoning goes awry for reasons not of form but of content.

When the fault lies in the form or structure of the argument, the fallacious inference is called a 'formal' fallacy. When it lies in the content of the argument, it is called an 'informal' fallacy. In the course of philosophical history, philosophers have been able to identify and name common types or species of fallacy. Oftentimes, therefore, the charge of fallacy calls upon one of these types.

Formal fallacies

We saw in 1.4 that one of the most interesting things about arguments is that their logical success or failure doesn't entirely depend upon their content, or what they claim. Validity is, again, content-blind or topic-neutral. The success of arguments in crucial ways depends upon how they *structure* their content. The following argument form is valid:

1. All Xs are Ys.
2. All Ys are Zs.
3. Therefore, all Xs are Zs.

For example:

1. All lions are cats. (true)
2. All cats are mammals. (true)
3. Therefore, all lions are mammals. (true)

With this form, whenever the premises are true, the conclusion must also be true (1.4). There's no way around it. With just a small change, however, in the way these Xs, Ys, and Zs are structured, validity evaporates, and the argument becomes invalid – which means, again, that it's no longer always the case that if the premises are true the conclusion must also be true.

1. All Xs are Ys.
2. All Zs are Ys.
3. Therefore, all Zs are Xs.

For example, substituting in the following terms results in true premises but a false conclusion.

1. All lions are cats. (true)
2. All tigers are cats. (true)
3. Therefore, all tigers are lions. (false)

This is an instance of showing *invalidity by counterexample* (1.5, 3.12). If this form were valid, it wouldn't be possible to assign content to it in a way that results in true premises but a false conclusion. The form simply wouldn't allow it. This is an important point. As we work our way through various fallacies in this book, pay attention to whether or not the fault in reasoning flows from a faulty form or something else.

Informal fallacies

What about fallacies that aren't rooted in a faulty form at all but instead in characteristically misleading content? How do they go wrong? A well-known example of an informal fallacy is the *gambler's fallacy* – it's both a dangerously persuasive and a hopelessly flawed species of inference.

The gambler's fallacy often occurs, for example, when someone takes a bet on the toss of a fair coin. The coin has landed heads up, say, seven times in a row. On the basis of this or a similar series of tosses, the fallacious gambler concludes that the next toss is more likely to come up tails than heads (or the reverse). What makes this an informal rather than a formal fallacy is that we can curiously present the reasoning here using a *valid form* of argument, even though the reasoning is bad.

1. If I've already tossed seven heads in a row, the probability that the eighth toss will yield a head is less than 50–50 (that is, a tails is due).
2. I've already tossed seven heads in a row.
3. Therefore, the probability that the next toss will yield a head is less than 50–50.

The *form* is perfectly valid; logicians call it *modus ponens*, the way of affirmation (see 3.1). Formally, *modus ponens* looks like this:

1. If *p*, then *q*.
2. *p*.
3. Therefore, *q*.

The flaw rendering the gambler's argument fallacious instead lies in the *content* of the first premise – the first premise is simply false. The probability of the next individual toss (like that of any individual toss) is and remains 50–50 no matter what toss or tosses preceded it.

Sure, the odds of tossing eight heads in a row are very low. But if seven heads in a row have already been tossed (a rare event, too), the chances of the sequence of eight in a row being completed (or broken) on the next toss is still just 50–50. Because this factual error about probabilities remains so common and so easy to commit, it has been classified as a fallacy and given a name. It's a fallacy, however, only in an informal way.

Now, logicians speak in these precise ways about fallacies (as 'formal' and 'informal'), but remember that sometimes ordinary speech deviates from logicians' technical usages. Sometimes any widely held though false belief is described as a 'fallacy'. Don't worry. As the philosopher Ludwig Wittgenstein (1889–1951) said, language is like a large city with lots of different avenues and neighbourhoods. It's alright to adopt different usages in different parts of the city. Just keep in mind where you are.

SEE ALSO

1.5 Invalidity
3.11 Conversion, contraposition, obversion
4.5 Conditional/biconditional

READING

★ S. Morris Engel (1974). *With Good Reason: An Introduction to Informal Fallacies*
★ Irving M. Copi (1986). *Informal Fallacies*
★ H. V. Hansen and R. C. Pinto (1995). *Fallacies: Classical and Contemporary Readings*
Scott G. Schreiber (2003). *Aristotle on False Reasoning*
★ Julian Baggini (2006). *The Duck that Won the Lottery*

1.8 Refutation

Samuel Johnson was not impressed by Bishop George Berkeley's argument that material substance does not exist. In his *Life of Johnson* (1791) James Boswell reported that, when discussing Berkeley's theory with him, Johnson once kicked a stone with some force and said, 'I refute it thus.'

READING

- ★ Plato (c.428–347 BCE). Dialogues *Meno, Euthyphro, Theaetetus*, and *Symposium*
- Richard Robinson (1950). *Definition*
- Ludwig Wittgenstein (2953). *Philosophical Investigations*, §43, §§65–66
- Nuel Belnap (1993). On rigorous definitions. *Philosophical Studies*, 72(2/3): 115–146

1.11 Certainty and probability

Seventeenth-century French philosopher René Descartes (1596–1650) is famous for claiming he had discovered the bedrock upon which to build a new science that could determine truths about the world with absolute certainty. The bedrock was an idea that could not be doubted, the *cogito* ('I think') – or, more expansively, as he put it in Part 1, §7 of his 1644 *Principles of Philosophy*, 'I think therefore I am' ('*cogito ergo sum*'). Descartes reasoned that it is impossible to doubt that you are thinking, for even if you're in error or being deceived or doubting, you are nevertheless thinking; and if you are thinking, you exist.

Ancient Stoics like Cleanthes (c.331–c.232 BCE) and Chrysippus (c.280–c.207 BCE) maintained that there are certain experiences of the physical and moral worlds that we simply cannot doubt – experiences they called 'cataleptic impressions'. Later philosophers like the eighteenth century's Thomas Reid (1710–96) believed that ordinary experience is improperly doubted and that God guarantees the veracity of our cognitive faculties. His contemporary, Giambattista Vico (1688–1744), reasoned that we can be certain about things artificial or human but not about the non-human, natural world. More recently, the Austrian philosopher Ludwig Wittgenstein (1889–1951) tried to show how it simply makes no sense to say that one doubts certain things. Some purported doubts (e.g. about whether the external world exists) are, according to Wittgenstein, meaningless.

Others have come to suspect that there may be little or nothing we can know with *certainty* and yet concede that we can still figure things out with some degree of *probability*. Hellenistic Academic sceptics such as Arcesilaus (c.240–c.315 BCE) and Carneades (214–c.129 BCE) seem to have argued for this view. Before, however, you go about claiming to have certainly or probably discovered philosophical truth, it will be a good idea to give some thought to what each concept means.

Types of certainty

Certainty is often defined as a kind of feeling or mental state (perhaps as a state in which the mind believes some X without any doubt at all). But defining certainty this way offers only a psychological account of the concept, and a psychological account fails to define when we are properly *warranted* in feeling this way. A more philosophical account of certainty would therefore add something about that sort of warrant – perhaps with the idea that a proposition may be properly accepted as certainly true when it is impossible for it to be false; alternatively, it may be properly accepted as certainly false when it is impossible for it to be true. Sometimes propositions that are certain in this way are called *necessarily true* and *necessarily false* (1.12).

The sceptical problem