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?'

Grounds for premises and Agrippa's trilemma?

There are several important accounts about how a premise can be acceptable. One is that the premise is itself the conclusion of a different, solid argument (perhaps a nested argument). As such, the truth of the premise has been demonstrated elsewhere. But it is clear that if this were the only kind of justification for the inclusion of a premise, we would face an infinite regress. That is to say, each premise would have to be justified by a different argument, the premises of which would have to be justified by yet another argument, the premises of which … *ad infinitum*.

Now, there are philosophers called *infinitists* for whom regresses of this sort are not problematic. Unless, however, one wishes to live with the infinite regress, one must find another way of determining sentences acceptable to serve as premises.

A compelling option for many has been to conceive of truths not as a hierarchy but rather as a network so that it's the case that justifications ultimately just circle back around to compose a coherent, mutually supporting but ultimately anchor-less web. The objective of philosophers and other theorists, from this point of view, becomes a project of conceptual weaving and embroidery, stitching together concepts and arguments in consistent and meaningful ways to construct a coherent conceptual fabric. Philosophers who conceive of truths, theories, and reasoning in this way are called *coherentists*.

Philosophers who object to infinite regresses of justification and who find in the coherentist vision just vicious circularity often look for something fundamental or foundational, a stopping point or bedrock for reasons and justification. Philosophers of this sort are often called *foundationalists*. There must be for foundationalists premises that stand in need of no further justification through other arguments. Let's call them 'basic premises'.

There's been a lot of ink spilled about what are to count as basic premises and why they are basic. By some accounts (called *contextualist*), the local context in which one is reasoning determines what's basic. For example, a basic premise might be, 'I exist'. In most contexts, this premise does not stand in need of justification. But if, of course, the argument is trying to demonstrate that I exist, my existence cannot be used as a premise. One cannot assume what one is trying to argue for.

Other kinds of philosophers have held that certain sentences are more or less basic for other reasons: because they are based upon self-evident or 'cataleptic' perceptions (*stoics*), because they are directly rooted in sense data (*positivists*), because they are grasped by a power called intuition or insight (*Platonists*), because they make up the framework of any possible inquiry and therefore cannot themselves be the objects of inquiry (*Kantians, Wittgensteinians*), because they are revealed to us by God (*theologians*), or because we grasp them using cognitive faculties certified by God (*Cartesians*).

Other philosophers, principally sceptics, have challenged the idea that an ultimate ground can be given at all for reasoning. Appeals to neither (1) regresses, nor (2) circles, nor (3) foundations ultimately work. The problem is an old one and has been popularly described as 'Agrippa's trilemma'. See Graeco-Roman Diogenes Laëritus's *Lives of Eminent Philosophers* (9.88–89) and Sextus Empiricus's *Outlines of Pyrrhonism* (PH 1.15.164) for the details.

Formally, then, the distinction between premises and conclusions is clear. But it is not enough to grasp this difference. In order to use these philosophical tools, one has to be able both to spot the explicit premises and to make explicit the unstated ones. The philosophical issues behind that distinction, however, are deep. Aside from the question of whether or not the conclusion follows from the premises, one must come to terms with the thornier questions related to what justifies the use of premises in the first place. Premises are the starting points of philosophical argument. One of the most important philosophical issues, therefore, must be the question of where and how one begins.

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 sun will rise tomorrow might be presented in a way that can easily be taken for a deductive argument:

1. The sun rises every day.
2. Tomorrow is a day.
3. Therefore, the sun will rise tomorrow.

Because of its similarity with deductive forms, one may be tempted to read the first premise as an 'all' sentence:

> The sun rises on *all* days (every 24-hour period) that there ever have been and ever will be.

The limitations of human experience, however (the fact that we can't experience every single day), justify us in forming only the less strong 'some' sentence:

> The sun has risen on every day (every 24-hour period) that humans have recorded their experience of such things.

This weaker formulation, of course, enters only the limited claim that the sun has risen on a small portion of the total number of days that have ever been and ever will be; it makes no claim at all about the rest.

But here's the catch. From this weaker 'some' sentence, one cannot construct a well-formed deductive argument of the kind that allows the conclusion to follow with the kind of certainty characteristic of deduction. In reasoning about matters of fact, one would like to reach conclusions with the certainty of deduction. Unfortunately, induction will not allow it. There's also another more complex problem lurking here that's perplexed philosophers: induction seems viciously circular. It seems in fact to assume the very thing it's trying to prove. Consider the following.

Assuming the uniformity of nature?

Put at its simplest, the problem of induction can be boiled down to the problem of justifying our belief in the uniformity of nature or even reality across space and time. If nature is uniform and regular in its behaviour, then what's been *observed* past and present (i.e. premises of an induction) is a sure guide to the so far *unobserved* past, present, and future (i.e. the conclusion of an induction).

The only basis, however, for believing that nature is uniform is the *observed* past and present. We can't then, it seems, go beyond observed events without assuming the very thing we need to prove – that is, that unobserved parts of the world operate in the same way as the parts we observe. In short, inductively proving that some bit of the world is like other bits requires already assuming that uniformities of that sort hold.

> *Induction undertakes to prove the world to be uniform in specific ways; but inductive inference already assumes that the world is relevantly uniform.*

We can infer inductively that the sun will rise tomorrow on the basis of what it's done in the past (i.e. that the future will resemble the past) only if we already assume that the future will resemble the past. Eighteenth-century Scot David Hume has remained an important philosopher in part precisely for his analysis of this problem.

Believing, therefore, that the sun may *possibly not* rise tomorrow is, strictly speaking, *not* illogical, since the conclusion that it must rise tomorrow does *not* inexorably follow from past observations.

A deeper complexity

Acknowledging the relative weakness of inductive inferences (compared to those of deduction), good reasoners qualify the conclusions reached through it by maintaining that they follow not with necessity but only *with probability* (i.e. it's just highly probably that the sun will rise tomorrow). But does this fully resolve the problem? Can even this weaker, more qualified formulation be justified? Can we, for example, really justify the claim that, on the basis of uniform and extensive past observation, it is *more probable than not* that the sun will rise tomorrow?

The problem is that there is no deductive argument to ground even this qualified claim. To deduce this conclusion successfully we would need the premise 'what has happened up until now *is more likely* to happen tomorrow'. But this premise is subject to just the same problem as the stronger claim that 'what has happened up until now *must* happen tomorrow'. Like its stronger counterpart, the weaker premise bases its claim about the future only on what has happened up until now, and such a basis can be justified only if we already accept the uniformity (or at least probable continuity) of nature. But again, the uniformity (or continuity) of nature is just what's in question.

A groundless ground?

Despite these problems, it seems that we can't do without inductive generalisations and inductive reasoning generally. They are (or at least have been so far!) simply too useful to refuse. Inductive generalisations compose the basis of much of our scientific rationality, and they allow us to think about matters concerning which deduction must remain silent. In short, we simply can't afford to reject the premise that 'what we have so far observed is our best guide to what is true of what we haven't observed', even though this premise cannot itself be justified without presuming itself.

There is, however, a price to pay. We must accept that engaging in inductive generalisation requires that we hold an indispensable belief which itself, however, must remain in an important way unjustified. As Hume puts it: 'All our experimental conclusions proceed upon the supposition that the future will be conformable to the past. To endeavour, therefore, the proof of this last supposition by probable arguments … must be evidently going in a circle, and taking that for granted, which is the very point in question' (*Enquiry Concerning Human Understanding*, 4.19). Can we accept reasoning and sciences that are ultimately groundless?

SEE ALSO

1.1 Arguments, premises, and conclusions
1.2 Deduction
1.7 Fallacies
2.4 Analogies
5.5 Hume's fork

READING

Francis Bacon (1620). *Novum Organum*
David Hume (1739). *A Treatise of Human Nature*, Bk 1, Part 3, Section 6

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' one another or to be 'contrary' to one another, or at least together imply contradiction or contrariety (3.10).

Statements are *contradictory* when they are opposite in 'truth value': when one is true the other is false, and vice versa. Statements are *contrary* when they can't both be true but, unlike contradictories, can both be false. With contraries, at least one is false.

Consistency, like contradiction and contrariety, are about comparing two or more different statements. A *single* sentence can, however, be *self-contradictory* when it makes an assertion that is necessarily false – often by conjoining two inconsistent sentences, such as *p* and not-*p* (1.12). You might call such a sentence self-inconsistent. (Compare this with the idea of the *paraconsistent* in 3.10.)

All this can be boiled down to a simple formulation: two or more statements are *consistent* when it's logically possible for them all to be true (a) in the same sense and (b) at the same time. Two or more statements are *inconsistent* when it is not possible for them all to be true in the same sense and at the same time.

Apparent and real inconsistency: the abortion example

At its most flagrant, inconsistency is obvious. If I say, 'All murder is wrong' and 'That particular murder was right', I am clearly being inconsistent, because the second assertion is clearly contrary to the first. (One might be false, both might be false, but both can't be true.) On a more general level, it would be a bald contradiction to assert both that 'all murder is wrong' and 'not all murder is wrong'. (One must be true and the other false.)

But sometimes inconsistency is difficult to determine. Apparent inconsistency may actually mask a deeper consistency – and vice versa.

Many people, for example, agree that it is wrong to kill innocent human beings. And many of those same people also agree that abortion is morally acceptable. One argument against abortion is based on the claim that these two beliefs are inconsistent. That is, critics claim that it is inconsistent to hold both that 'It is wrong to kill innocent human beings' and that 'It is permissible to destroy living human embryos and fetuses'.

Defenders of the permissibility of abortion, on the other hand, may retort that properly understood the two claims are not inconsistent. A defender of abortion could, for example, claim that embryos are not human beings in the sense normally understood in the prohibition (e.g. conscious or independently living or already-born human beings). The defender, in other words, might return a rejoinder to the critic that her objection is based on an equivocation (3.3). Alternatively, a defender of abortion might modify the prohibition itself to make the point more clearly (e.g. by claiming that it's wrong only to kill innocent human beings that have reached a certain level of development, consciousness, or feeling).

Exceptions to the rule?

But is inconsistency always undesirable? Some people are tempted to say it is not. To support their case, they present examples of statements that intuitively seem perfectly acceptable yet seem to meet the definition of inconsistency. One example might be:

> It is raining, and it is not raining.

Of course, the inconsistency might be only apparent. What one may actually be saying is not that it's raining and not raining, but rather that it's neither properly raining nor not raining, since there is a third possibility – perhaps that it is drizzling, or intermittently raining – and that this other, *fuzzy* possibility most accurately describes the current situation (3.1).

What makes the inconsistency only apparent in this example is that the speaker is shifting the sense of the terms being employed. Another way of saying the first sentence, then, is that, 'In one sense it is raining, but in another sense of the word it is not'. For the clauses composing this sentence to be truly inconsistent, the relevant terms being used must retain precisely the same meaning throughout. But, when you do unearth a genuine logical inconsistency, you've accomplished a lot, for it can be very difficult if not impossible to defend the inconsistency without rejecting rationality outright. There are poetic, religious, and philosophical contexts, however, in which this is precisely what people find it proper to do.

Poetic, religious, or philosophical inconsistency?

The Danish existentialist philosopher Søren Kierkegaard (1813–55) maintained that the Christian notion of the incarnation ('Jesus is God, and Jesus was a man') is a paradox, a contradiction, an affront to reason, but nevertheless true (7.6). Many Christians simply hold the idea to be a difficult mystery.

That kind of difficulty, however, may extend farther than religious contexts. Atheist existentialist philosopher Albert Camus (1913–60) maintained that there is something fundamentally 'absurd' (perhaps inconsistent?) about human existence. Post-structuralist philosopher Jacques Derrida's theory of *différance* raises metaphysical questions about the consistency of reality (6.2). Philosophical fiction and poetry may enlist rhetorical strategies involving inconsistency (7.4). Dialetheists and others have even challenged the idea that consistency is fundamental to logic (3.10). Perhaps, then, Emerson was right, and there are contexts in which inconsistency and absurdity paradoxically make sense.

Consistency ≠ truth

Be this as it may, inconsistency in philosophy is generally a serious vice. Does it follow from this that consistency is philosophy's highest virtue? Not quite. Consistency is only a minimal condition of acceptability for a philosophical position. Since it's often the case that one can hold a consistent theory that is inconsistent with another, equally consistent theory, the internal consistency of any particular theory is no guarantee of its truth. Indeed, as French philosopher-physicist Pierre Maurice Marie Duhem (1861–1916) and the American philosopher Willard Van Orman Quine (1908–2000) have separately maintained, it may be possible to develop two or more theories that are (1) internally consistent, yet (2) inconsistent with each other, and also (3) perfectly consistent with all the data we can possibly muster to determine the truth or falsehood of the theories (7.11).

Take as an example the so-called problem of evil. How do we solve the puzzle that God is supposed to be good but that there is also awful suffering (an apparent evil) in the world? As it turns out, you can advance a number of theories that may solve the puzzle but remain inconsistent with one another. You can hold, for instance, that God does not exist. Or you can hold that God allows suffering for a greater good. Although each solution may be perfectly consistent with itself, they can't both be right, as they are inconsistent with each other. One theory asserts God's existence, and the other denies it. Establishing the consistency of a position, therefore, may advance and clarify philosophical thought, but it probably won't settle the issue at hand. We often need to appeal to more than consistency if we are to decide between competing positions. How we do this is a complex and controversial subject of its own.

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.'

Any great person is allowed one moment of idiocy to go public, and Johnson's attempt at a refutation must be counted as just such a moment, because he wildly missed Berkeley's point. The bishop would never have denied that one could kick a stone; he denied that stones properly understood can be conceived to be material substances. But Johnson's refutation also failed even to be a true *refutation*, a concept that in philosophy has a precise meaning.

To refute an *argument* is to show that its reasoning is bad. If you, however, merely register your disagreement with an argument, you are not refuting it – even though in everyday speech people often talk about refuting a claim in just this way. So, how can one really refute an argument?

Refutation tools

There are two basic ways of doing this, both of which are covered in more detail elsewhere in this book. First, you can show that the argument is invalid: the conclusion does not follow from the premises as claimed (see 1.5). Or, second, you can show that one or more of the premises are false (see 1.4).

There is a third method of refutation, too – or at least quasi-refutation. All you have to do is simply show that the conclusion *must be false*. From this it can be argued that therefore, even if you can't identify exactly what is wrong with the argument, *something must be wrong* with it (see 3.25). This last method, however, isn't strictly speaking a refutation, since one has failed to show *what* is wrong with the argument, only *that* it must be wrong. Nevertheless, this understanding that something must be wrong often accomplishes all that's needed.

Inadequate justification

Refutations are powerful tools, but it would be rash to conclude that in order to reject an argument *only* a refutation will do. You may be justified in rejecting an argument even if you have not strictly speaking refuted it. You may not be able to show that a key premise is false, for example, but you may believe that it's inadequately justified. An argument based on the premise that 'there is intelligent life elsewhere in our universe' would fit this model. We can't show that the premise is actually false, but we can argue that we have both no good reasons for believing it to be true and some grounds for supposing it to be false. Therefore, we can regard any argument that depends on this premise as rather dubious and permissibly ignore it.

Conceptual problems

More contentiously, you might also reject an argument by arguing that it utilises a concept inappropriately. This sort of problem is particularly clear in cases where a vague concept is used as if it were precise. For instance, consider the claim that the government is obliged to provide assistance only to those who do not have enough to live on properly. But given that there can be no precise formulation of what 'enough to live on properly' means, any argument must be inadequate that concludes by making a sharp distinction between those who have enough in this sense and those who don't. The logic of the argument may be impeccable and the premises may appear to be true. But if you use vague concepts in precise arguments you may well end up with distortions.

Using the tool

There are many more ways of legitimately objecting to an argument without actually refuting it. The important thing is to keep in mind the difference between refutation and other forms of objection and to be clear about what form of objection you're offering.

SEE ALSO

1.4 Validity and soundness
1.5 Invalidity
3.4 Bivalence and the excluded middle

READING

Imre Lakatos (1976/2015). *Proofs and Refutations*
Karl Popper (1963). *Conjectures and Refutations*
★ Jamie Whyte (2005). *Crimes Against Logic*

★ Julian Baggini (2008). *The Duck That Won the Lottery and 99 Other Bad Arguments*
★ T. Schick, Jr, and L. Vaughn (2020). *How to Think about Weird Things*, 8th edn

1.9 Axioms

Obtaining a guaranteed true conclusion in a deductive argument requires that the argument be *sound* – that is, it requires both (1) that the argument be valid and (2) that the premises be true (1.4). Unfortunately, the procedure for deciding whether or not a premise is true is much less determinate than the procedure for assessing an argument's validity. Unless premises are to be justified by arguments whose own premises are to be justified by still other arguments *ad infinitum*, and unless premises are to circle back on themselves in a loop of justification, there must be a stopping point where fundamental or basic premises are just accepted as true (see Agrippa's trilemma in 1.1).

Defining axioms

For this reason, the concept of an *axiom* becomes a useful philosophical tool. An axiom is a proposition that acts as a special kind of premise in a specific kind of rational system. Axiomatic systems were first formalised by the Alexandrian geometer Euclid (fl. 300 BCE) in his famous work the *Elements*. In these kinds of systems, axioms function as initial, anchoring claims that stand in no need of justification – at least from within the system. They are the bedrock of the theoretical system, the basis from which, through various steps of deductive reasoning, the rest of the system is derived. In ideal circumstances, an axiom should be such that no rational agent could possibly object to its use.

Axiomatic vs natural systems of deduction

It is important to understand, however, that not all conceptual systems are axiomatic – not even all rational systems. For example, some deductive systems try simply to replicate and refine the procedures of reasoning that seem to have unreflectively or naturally developed among humans. This type of system is called a *natural system of deduction;* it doesn't posit any axioms but looks instead for its formulae to the practices of ordinary rationality.

First type of axiom

As we have defined them, axioms would seem to be pretty powerful premises. Once, however, you consider the types of axiom that there are, their power seems to be somewhat diminished. One type of axiom comprises premises that are true by definition. Perhaps because so few great philosophers have been married, the example of 'all bachelors are unmarried men' is usually offered as the paradigmatic example of definitional truths. The problem is that no argument is going to be able to run very far with such an axiom. Axioms of this sort are purely tautological, that is to say, 'unmarried men' merely restates in different words the meaning that is already contained in 'bachelor'. (This sort of proposition is sometimes called – following Immanuel Kant – an 'analytic' proposition. See 4.3.) They are therefore spectacularly uninformative sentences (except to someone who doesn't know what 'bachelor' means). So, they are unlikely to help yield informative conclusions in an argument.

Second type of axiom

Another type of axiom is also true by definition, but in a slightly more interesting way. Many regions of mathematics and geometry rest on their axioms, and it's only by accepting these basic axioms that more complex proofs can be constructed within those regions. You might call these propositions 'primitive' sentences within the system (7.7). For example, it is an axiom of Euclidean geometry that the shortest distance between any two points is a straight line. But while axioms like these are vital in geometry and mathematics, they merely *stipulate* what is true *within* the particular system of geometry or mathematics to which they belong. Their truth is guaranteed, but only in a limited way – that is, only *within the context* in which they're defined. Used in this way, axioms' acceptability rises or falls with the acceptability of the theoretical system as a whole.

Axioms for all?

So, some axioms aren't terribly informative, while others are limited to specific contexts. Some may find this account rather unsatisfactory and object to it. Aren't there any 'universal axioms' that are both secure and informative in all contexts universally, for all thinkers, no matter what? Some philosophers have thought so.

The Dutch philosopher Baruch (also known as Benedictus) Spinoza (1632–77) in his *Ethics* (1677) attempted to construct an entire metaphysical system from just a few axioms, axioms that he believed to be universal truths virtually identical with God's thoughts. The problem is that most would agree that at least some of his axioms seem to be empty, unjustifiable, and parochial assumptions. For example, one of Spinoza's axioms states that 'if there be no determinate cause it is impossible that an effect should follow' (*Ethics*, Bk 1, Pt 1, axiom 3).

As English empiricist John Locke (1632–1704) argues, however, this claim, taken literally, is pretty uninformative, since it's true by definition that all effects have causes. What the axiom seems to imply, however, is a more metaphysical claim – that all events in the world are effects that necessarily follow from their causes. Working in Locke's wake, David Hume (1711–76) points out that the metaphysical claim fares no better. Not only do we have no reason to think it's true, but moreover it's not at all senseless to hold that an event might occur without any cause at all (*Treatise*, 1.3.14). Medieval Islamic philosopher al-Ghazali (1058–1111) advanced a similar line in his *The Incoherence of the Philosophers* ('On natural science', Question 1ff.).

Of course, Spinoza seems to claim that he has grasped the truth of his axioms through a special form of intuition (*scientia intuitiva*), and many philosophers have held that there are 'basic' and 'self-evident' truths that may serve as axioms in our reasoning. (See 7.1.) But why should we believe them?

In many contexts of rationality, therefore, axioms seem to be a useful device, and axiomatic systems of rationality often serve us very well, indeed – especially as part of mathematics and logical theory. But the notion that those axioms can be so secure that no rational person could in any context deny them seems to be rather dubious.

SEE ALSO

3.6 Circularity
4.6 Cause/reason

7.1 Basic beliefs
7.8 Self-evident truths

READING

Euclid (c. 300 BCE). *Elements*
★ Alfred Tarski (1946/95). *Introduction to Logic and to the Methodology of Deductive Sciences*
A.A. Fraenkel, Y. Bar-Hillel, and A. Levy (1973). *Foundations of Set Theory*
Fred R. Berger (1977). *Studying Deductive Logic*

1.10 Definitions

If, somewhere, there lie written on tablets of stone the ten philosophical commandments, you can be sure that numbered among them is the injunction to 'define your terms'. In fact, definitions are so important in philosophy that some have maintained that definitions are ultimately all there is to the subject.

Definitions are important because without them, it's very easy to argue at cross-purposes or to commit fallacies involving equivocation (3.3). As the experience of attorneys who questioned former US president Bill Clinton show, if you are, for example, to interrogate someone about extramarital sex, you need to define what precisely you mean by 'sex'. Otherwise, much argument down the line, you can bet someone will turn around and say, 'Oh, well, I wasn't counting *that* as sex'. Much of our language is vague and ambiguous, but if we are to discuss matters in as precise a way as possible, as philosophy aims to do, we should remove as much vagueness and ambiguity as possible, and adequate definitions are the perfect tool for helping us do that.

Free trade example

For example, consider the justice of 'free trade'. In doing so, you may define free trade as 'commercial exchange that is not hindered by national or international law'. But note that with this rendering you have fixed the definition of free trade for the purposes of your discussion. Others may argue that they have better or alternative definitions of free trade. This may lead them to reach different conclusions about its justice. You might respond by adopting a new definition, defending your original definition, or proposing yet another definition. And so it goes. That's why setting out definitions for difficult concepts and reflecting on their implications composes a great deal of philosophical work.

Again, the reason why it's important to lay out clear definitions for difficult or contentious concepts is that any conclusions you reach properly apply only to those concepts (e.g. 'free trade') *as defined*. A clear definition of how you will use the term thereby both helps and constrains discussion. It helps discussion because it gives a determinate and non-ambiguous meaning to the term. It limits discussion because it means that whatever you conclude does not necessarily apply to other uses of the term. As it turns out, much disagreement in life results from the disagreeing parties, without their realising it, meaning different things by their terms.

Too narrow or too broad?

That's why it's important to find a definition that does the right kind of work. If one's definition is *too narrow* or idiosyncratic, it may be that one's findings cannot be applied as broadly as could be hoped. For example, if one defines 'man' to mean bearded, human, male adult, one may reach some rather absurd conclusions – for example, that many Native American males are not men. A tool for criticism results from understanding this problem. In order to show that a philosophical position's use of terms is inadequate because *too narrow*, point to a case that ought to be covered by the definitions it uses but clearly isn't.

If, on the other hand, a definition is *too broad*, it may lead to equally erroneous or misleading conclusions. For example, if you define wrongdoing as 'inflicting suffering or pain upon another person' you would have to count the administering of shots by physicians, the punishment of children and criminals, and the coaching of athletes as instances of wrongdoing. Another way, then, of criticising someone's position on some philosophical topic is to indicate a case that fits the definition he or she is using but which should clearly not be included under it. Cases showing that definitions are too broad are special kinds of *counterexample* (3.12).

A definition is like a property line; it establishes the limits marking or defining those instances to which it's proper to apply a term and those instances to which it is not. In this sense, a definition articulates the *specific differences* that distinguish one kind of thing from all others (5.2). The ideal definition, therefore, permits application of the term to just those cases to which it should apply – and to no others. It will admit no counterexamples.

Often, philosophers attempt to figure relatively perfect definitions by thinking through both the *sufficient and necessary conditions* for using a concept or term. Elaborating (perhaps not terribly well) on Aristotle's famous definition, one might formulate the sufficient and necessary conditions for being a human by saying that something is 'human' *if and only if* it's a rational, risible, fine-haired, bipedal primate (see 4.17). Another way to think of a definition is as a special kind of *definite description*, a formulation that well describes what it defines (4.14).

A rule of thumb

As a general rule, it's better if your definition corresponds as closely as possible to the way in which the term is ordinarily used in the kinds of debates to which your claims are pertinent. There will be, however, occasions where it is appropriate, even necessary, to coin *special uses* through what philosophers call *stimulative definition*. This would be the case where the current lexicon is not able to make distinctions that you think are philosophically important. For example, we do not have a term in ordinary language that describes a memory that is not necessarily a memory of something the person having it has experienced. Such a thing would occur, for example, if I could somehow share your memories: I would have a memory-type experience, but this would not be of something that I had actually experienced. To call this a memory would be misleading. For this reason, philosophers have coined the special term 'quasi-memory' (or 'q-memory') to refer to these hypothetical memory-like experiences.

A long tradition

Historically, many philosophical questions are, in effect, quests for adequate definitions. What is knowledge? What is beauty? What is the good? Here, it's not enough just to say, 'By knowledge I usually mean something like …'. Rather, the search is for a definition that *best* articulates the concept in question and does so in as general or universal a way as possible. Much of the philosophical work related to definition takes the form of *conceptual analysis* or the attempt to unpack and clarify the meanings of important concepts. What is to count as the best articulation or a proper analysis, however, requires a great deal of debate. Indeed, it's a viable philosophical question as to whether or not many philosophy concepts actually can be defined. Perhaps some concepts are so complex that they can't be compressed into a reasonably compact formulation. Perhaps the best that can be done is to become familiar with their usages by just diving into the network of philosophical theory in which they appear.

Many philosophers have not been deterred. For some that's because of their philosophical commitments concerning the nature of reality and human epistemic powers. Ancient and medieval thinkers (like Plato and Aquinas), for example, seem to have been confident about the project of formulating adequate definitions because they were committed to the idea that reality includes *essences* or *natures* that exist independently of us and that define what things truly are (4.12). Moreover, these thinkers were convinced that human beings possess the capacity to apprehend those essences and formulate them in language. Many more recent thinkers (like some pragmatists and post-structuralists) have held that definitions are nothing more than conceptual instruments that organise our interactions with each other and the world. That is so because recent philosophy has in large measure abandoned the idea that human language can meaningfully formulate real, independent essences or even that such essences exist.

The labor of analysing concepts has been related too to philosophical criticisms of philosophy itself. Some thinkers have gone so far as to argue that virtually all philosophical problems are at the end of the day rooted in nothing more than failures to understand how ordinary language functions. Resolving those puzzles, from this point of view, entails clarifying the way we use language so as to eliminate the confusions upon which philosophy generates its conundrums. While, to be accurate, this project demands more than just scrutinising definitions, it does show just how deep the philosophical preoccupation with getting language right runs.

SEE ALSO

3.9 Criteria
4.14 Knowledge by acquaintance/description
4.17 Necessary/sufficient
5.9 Signs and signifiers

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

The main problem, philosophically speaking, thinkers face is in establishing that it is in fact impossible for any candidate for certainty to have a different truth value. Sceptical thinkers have been extremely skillful in showing how virtually any claim might possibly be false even though it appears to be true (or possibly true though it appears to be false). In the wake of sceptical scrutiny, many agree that absolute certainty in advancing truth claims remains unattainable. One reason for this is the question of whether or not one must be certain that one is certain. (Can you be sure that you're really sure?)

These are serious though perhaps not insurmountable problems for certainty. For many they present deep sceptical trouble for anyone interested in apprehending truth. On the other hand, clearly not all that's true is certain. So, perhaps certainty isn't required for making truth claims or claims to having acquired knowledge. Is there a way to leave the problems of certainty behind and still confidently determine uncertain truths? What is the next best thing if we give up on certainty? To give a proper answer to this question would require a much larger study of *epistemology* or the theory of knowledge. But for the sake of our concerns here, consider the answer that's most commonly advanced: *probability*.

Probability is the natural place to retreat to if certainty becomes intolerably problematic. What is merely probable also seems the largest fraction of human epistemic life. As John Locke writes in his 1689 *Essay Concerning Human Understanding*: 'the greatest part of our concernments' are 'only the twilight, as I may so say, of probability' (4.14.2). As a refuge, however, probability is rather like the house of sticks to which the little pig flees when the wolf arrives at the door of his straw house. Probability faces vulnerabilities of its own.

Objective and subjective probability

We can distinguish between objective and subjective probability. *Objective* probability is where what will happen is genuinely indeterminate. Radioactive decay could be one example. For any given atom of a radioactive material, the probability of it decaying over the period of its half-life is 50–50. This means that, if you were to take ten such atoms, it is likely that five will decay over the period of the element's half-life, while five will not decay. On at least some interpretations in physics, it's genuinely indeterminate which atoms will fall into which category.

Subjective probability, on the other hand, refers to cases where there may be no actual indeterminacy, but some particular mind or set of minds makes a probability judgement about the likelihood of some event. These subjects do so because they lack complete information about the causes that will determine the event. Their ignorance requires them to make a probabilistic assessment, usually by assigning a probability based on the number of occurrences of each outcome over a long sequence in the past.

So, for example, if we toss a coin, cover it, and ask you to bet on heads or tails, the outcome has already been determined. Since you don't know what it is, you have to use your knowledge that heads and tails over the long run fall 50–50 to assign a 50 per cent probability that it's a head and a 50 per cent probability that it's a tail. If you could see the coin, there would be no 50–50 about it. You'd know the side that's up with, in fact, 100 per cent certainty.

The odds set by gamblers and handicappers at horse races are also species of subjective probability. The posted odds record simply what the many people betting on the race subjectively believe about the outcome, not the real chance of any horse's crossing the finish line first.

Certainty and validity

If you have a valid deductive argument, then its conclusion is often said to follow from the premises with certainty. Many inquirers, however, demand not only that conclusions *follow* with certainty but that the conclusions themselves be certainly true. Consider the difference between the following arguments:

1. If it rained last night, England will probably win the match.
2. It rained last night.
3. Therefore, England will probably win the match.

1. It's certainly true that no parallel lines intersect.
2. These two lines are parallel.
3. Therefore, these two lines certainly do not intersect.

The conclusion of the first argument clearly enters only a probable claim. The conclusion of the second argument, in contrast to the first, enters a certain claim. But here's the rub: both examples present valid deductive arguments. Both arguments possess valid forms. Therefore, in both arguments the conclusion *follows* with certainty – i.e. the truth of the premises *guarantees* the truth of the conclusion – even though the *content* of one conclusion enters merely a probable claim, while that of the other enters a claim of certainty.

You must therefore distinguish between (1) whether or not the conclusion of an argument *follows* from the premises with certainty or some probability, and (2) whether or not the conclusion of an argument advances a *statement* the *content* of which concerns matters of probability or certainty.

Philosophical theories

But what about philosophical theories? It would seem that if certainty in philosophical theories were attainable, there would be little or no dispute among competent philosophers about which are true and which false – but, in fact, there seems to be a lot of dispute. Does this mean that the truth of philosophical theories is essentially indeterminate? Is deep disagreement a fundamental characteristic of philosophical inquiry?

Some philosophers would say no. For example, they would say that although there remains a great deal of dispute, there is also near unanimous agreement among philosophers on many things – for example, that Plato's theory of metaphysical forms is false and that Cartesian mind–body dualism is untenable.

Others of a more sceptical bent are, if you'll pardon the pun, not so certain about the extent to which anything has been proven, at least with certainty, in philosophy. Accepting a lack of certainty can from their point of view be seen as a matter of philosophical maturity.

SEE ALSO

1.2 Deduction
1.4 Validity and soundness
1.5 Invalidity
1.9 Axioms
1.12 Tautologies, self-contradictions, and the law of non-contradiction

READING

Ludwig Wittgenstein (1969). *On Certainty*, §115, §341

★ Barbara J. Shapiro (1983). *Probability and Certainty in Seventeenth-Century England*

Peter Klein (1992). Certainty. In: *A Companion to Epistemology* (eds J. Dancy and E. Sosa), 61–64

★ D.H. Mellor (2005). *Probability: A Philosophical Introduction*

★ Alan Hájek (2019). Interpretations of probability. In: *The Stanford Encyclopedia of Philosophy* (ed. Edward N. Zalta), Fall 2019 edn

1.12 Tautologies, self-contradictions, and the law of non-contradiction

Tautology and self-contradiction fall at opposite ends of a spectrum: the former is a sentence that's necessarily true, and the latter a sentence that's necessarily false. Despite being in this sense poles apart, they're actually intimately related.

In common parlance, *tautology* is a pejorative term used to deride a claim because it purports to be informative but in fact simply repeats the meaning of something already understood. For example, consider: 'The criminal has broken the law'. This statement might be mocked as a tautology since it tells us nothing about the criminal to say he has broken the law. To be a lawbreaker is precisely what it means to be a criminal.

In logic, however, 'tautology' has a more precisely defined meaning. A tautology is a statement that, because of its logical structure, is true in every circumstance – or, as some say, in every possible world. Tautologies are in this sense *logical truths* or *necessary truths*. Take, for example:

p or not-*p*.

If *p* is true the statement turns out to be true. But if *p* is false, the statement still turns out to be true. This is the case for *whatever* one substitutes for *p*: 'today is Monday', 'atoms are invisible', or 'monkeys make great lasagna'. One can see why tautologies are so poorly regarded. A statement that is true regardless of the truth or falsehood of its components can be considered to be empty; its content does no work.

This is not to say that tautologies are without philosophical value. Understanding tautologies helps one to understand the nature and function of reason and language.

Valid arguments as tautologies

As it turns out, all valid arguments can be restated as tautologies – that is, hypothetical statements in which the antecedent is the conjunction of the premises and the consequent the conclusion. In other words, every valid argument may be articulated as a statement of this form: 'If *w*, *x*, and *y* are true, then *c* is true', where *w*, *x*, and *y* are the argument's premises and *c* is its conclusion. When any valid argument is substituted into this form, a tautology results.

Law of non-contradiction

In addition, the law of non-contradiction – a cornerstone of philosophical logic – is also a tautology. The law may be formulated this way:

Not (*p* and not-*p*).

The law is a tautology since, whether *p* is true or false, the complete statement will turn out to be true.

The law of non-contradiction can hardly be said to be uninformative, since it's the foundation upon which nearly all logic is built. But, in fact, it's not the law itself that's informative so much as any attempt to break it.

Attempts to break the law of non-contradiction themselves require contradictions, and it's standardly accepted that contradictions are obviously, and in all circumstances, false. A contradiction flouts the law of non-contradiction, since it asserts both that something is true and that something is false in precisely the same sense and at the same time – asserting, as it were, both *p* and not-*p*. Given, however, that the law of non-contradiction is a tautology, and thus in all circumstances true, there can be nothing more clearly flawed and senseless than asserting a contradiction in opposition to it – unless, that is, you're a dialetheist in logic (see 3.10).

The principle of non-contradiction has also been historically important in philosophy. The principle underwrote ancient analyses of change and plurality and is crucial to Parmenides of Elea's sixth-century BCE proclamation that 'what-is is and cannot not-be'. It also seems central to considerations of identity – for example, in Leibniz's claim that objects that are identical must have all the same properties.

Self-refuting criticism

One curious and useful feature of the law of non-contradiction is that, as Aristotle shows in his *Metaphysics* Book 4, any attempt to refute it presupposes it, and so for Aristotle nothing can be more certain than the principle of non-contradiction. (See also Plato's formulation at *Republic* IV, 436b–437a.)

To argue that the law of non-contradiction is false is to imply that it is not also true. In other words, the critic *presupposes* that what he or she is criticising can be *either* true or false *but not both true and false*. But this presupposition is just the law of non-contradiction itself – the same law the critic aims to refute! In other words, anyone who denies the principle of non-contradiction simultaneously affirms it. It is, in short, a principle that cannot be rationally criticised, because it's presupposed by all rationality.

To understand why a *tautology* is necessarily, and in a sense at least, uninformatively true, and why a *self-contradiction* is necessarily false, is to understand the most basic principle of logic. The *law of non-contradiction* is where those two concepts meet and so is perhaps best described as the keystone, rather than cornerstone, of philosophical logic.

SEE ALSO

1.4 Validity and soundness
1.6 Consistency
3.10 Contradiction/contrariety
5.6 Leibniz's law of identity
7.5 Paradoxes

READING

Aristotle (384–322 BCE). *Interpretation*, esp. Chs 6–9
Aristotle (384–322 BCE). *Posterior Analytics*, Bk 1, Ch. 11:10
Graham Priest, J.C. Beall, and Bradley Armour-Garb (eds) (2004). *The Law of Non-contradiction* (2004)