chapter 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 are concerned with the way in which beliefs we have about the world either are or are not supported 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 arguments only if one scrutinizes their content and structure with supreme diligence.

Argument.

What, then, is an argument? 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 – but need not – include reasoning.

Philosophers, by 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 reason. An 'argument' 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').

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; 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 to 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 is not universally realized 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. Philosophers do indeed advance such ideas and theories, but in most cases their power and scope stems 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 distinguishing philosophy from them. (In fact, whether or not it is possible to do so 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 consciousness studies, theoretical physics, and anthropology). But theories composing these sciences are generally determined through certain formal procedures of experimentation and reflection with which philosophy has little truck. 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 religious 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). Furthermore, many wish to 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.

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

Philosophy, then, is not the only field of thought for which rationality is important. And not all that goes by the name of philosophy may be 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, literally, that with which an argument concludes, the product and result of a chain of inference, that which the reasoning justifies and supports.

What about 'premises'? In the first place, in order for a sentence to serve as a premise, it must exhibit this essential property: it must make a claim that is either true or false. Sentences do many things in our languages, and not all of them have that property. Sentences that issue commands, for example, ('Forward march, soldier!'), or ask questions ('Is this the road to Edinburgh?'), or register exclamations ('Holy cow!'), are neither true nor false. Hence it is not possible for them to serve as premises.

This much is pretty easy. But things can get sticky in a number of ways.

One of the most vexing issues concerning premises is the problem of implicit claims. That is, in many arguments key premises 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'.

In working out precisely what the premises are in a given argument, ask yourself first what the 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.

Indicators.

Sometimes certain words and phrases will indicate premises and conclusions. Phrases like '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', often indicate premises. (For example, 'Since the DNA, the fingerprints and the eyewitness accounts all implicate Smithers, she must be the killer.')

Premises, then, compose the set of claims from which the conclusion is drawn. In other sections, the question of how we can justify the move from premises to conclusion will be addressed (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?

There are two basic reasons why a premise might be acceptable. One is that the premise is itself the conclusion of a different, solid 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.* (In fact, sceptics – Eastern and Western, modern and ancient – have pointed to just this problem with reasoning.)

So unless one wishes to live with the problem of the infinite regress, there must be another way of finding sentences acceptable to serve as premises. There must be, in short, premises that stand in need of no further justification through other arguments. Such premises may be true by definition. (An example of such a premise is 'all bachelors are unmarried'.) But the kind of premises we're looking for might also include premises that, though conceivably false, must be taken to be true for there to be any rational dialogue at all. Let's call them 'basic premises'.

Which sentences are to count as basic premises depends on the context in which one is reasoning. One example of 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.

Philosophers have held that certain sentences are more or less basic for various 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 are revealed to us by God (Jewish, Christian and Islamic philosophers), or because we grasp them using cognitive faculties certified by God (Descartes, Reid, Plantinga). In our view, a host of reasons, best described as 'context' will determine them.

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 to spot the explicit premises and make explicit the unstated ones. And aside from the question of whether or not the conclusion follows from the premises, one must come to terms with the thornier question of what justifies the use of premises in the first place. Premises are the starting points of philosophical argument. As in any edifice, intellectual or otherwise, the construction will only stand if the foundations are secure.

See also

1.2 Deduction
1.3 Induction
1.9 Axioms
1.10 Definitions
3.6 Circularity
6.1 Basic beliefs
6.6 Self-evident truths

Reading

*Nigel Warburton, *Thinking From A to Z*, 2nd edn (2000)
*Patrick J. Hurley, *A Concise Introduction to Logic*, 7th edn (2000)

1.2 Deduction

> The murder was clearly premeditated. 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* 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 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 these people?

You might well be thinking there is 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, I believe, not true, because one or both (in this case both) premises are not true. 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 given: a (successful) deductive argument is one where, if the premises are true, then the conclusion must also be true.

But before we leave this topic, we should return to the investigations of our detective. Reading his deliberations, one could easily insert the vital, missing word. 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 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 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 ones is much trickier. To judge whether the conclusion really *must* follow from the premises, we 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 it. 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

*John Shand, *Arguing Well* (2000)
Fred R. Berger, *Studying Deductive Logic* (1977)

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 but only with probability (though even this formulation is problematic, as we will see).

Defining induction.

Often, induction involves reasoning from a limited number of observations to wider, probable generalizations. Reasoning this way is commonly called 'inductive generalization'. It is 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 so far as human experience can recall, so people reason that it will probably rise tomorrow. (The work of the Scottish philosopher David Hume [1711–76] has been influential on this score.) 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 m 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 problem of induction.

Inductive generalizations are, however, often where the action is. Reasoning in experimental science, for example, depends on them in so far as scientists formulate and confirm universal natural laws (e.g. Boyle's ideal gas law) on the basis of a relatively small number of observations. The tricky thing to keep in mind about inductive generalizations, however, is that they involve reasoning from a 'some' in a way that only works with *necessity* for an 'all'. This type of inference makes inductive generalization 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 (Example A):

1. *Some* 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 be true and the conclusion still be false. Properly understood, however, it may be a strong inductive argument – for example, if by ‘some’ elephants one means ‘all but one’ and if the conclusion is interpreted to mean ‘it is *probably* the case that this elephant likes chocolate’.

On the other hand, consider this rather similar argument (Example B):

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. The problem of induction is 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 kind of 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 might 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.

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. If nature is uniform and regular in its behaviour, then events in the *observed* past and present are a sure guide to unobserved events in the *unobserved* past, present and future. But the only grounds for believing that nature is uniform are the *observed* events in the past and present. We can't seem to go beyond the events we observe 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've observed. (This is just the problem to which Hume points.) 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 *with probability*. 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* that the sun will rise tomorrow than it won't?

Strictly speaking 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 *is certain* to 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 accept the uniformity (or at least general continuity) of nature. But 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 generalizations. They are (or at least have been so far!) simply too useful to refuse. They compose the basis of much of our scientific rationality, and they allow us to think about matters concerning which deduction must remain silent. 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 by deductive argument.

There is, however, a price to pay. We must accept that engaging in inductive generalization requires that we hold an indispensable belief which itself, however, must remain in an important way ungrounded.

See also

1.1 Arguments, premises and conclusions
1.2 Deduction
1.7 Fallacies
2.4 Analogies
3.14 Hume's Fork

Reading

*David Hume, *A Treatise of Human Nature* (1739–40), bk 1

1.4 Validity and soundness

In his book *The Unnatural Nature of Science* the eminent British biologist Lewis Wolpert 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 sciences on this point. Its theories, conclusions and terms can at times be extraordinarily counter-intuitive 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 it is a perfect example of a valid argument. What's going on?

Defining validity.

Validity is a property of well-formed deductive arguments, which, to recap, is defined as an argument where the conclusion is in some sense (actually, hypothetically, etc.) presented as following from the premises *necessarily* (see 1.2). A valid deductive argument is one for which the conclusion follows from the premises in that way.

The tricky thing, however, 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*. And so, with regard to validity, the *content* or truth of the statements composing the argument is 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 one doesn't need to attach any meaning to the variables to see that this particular structure is a valid one. No matter what we replace the variables with, it will always be the case that *if* the premises are true (although 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, then it is an invalid argument.

What this boils down to is that the notion of validity is 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 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.

From another point of view we might consider that deductive arguments work a bit like sausage machines. You put ingredients (premises) in, and then you get something (conclusions) out. Deductive arguments are the best kind of sausage machine because they *guarantee* that when you put good ingredients (all true premises) in, you get a quality product (true conclusions) out.

A good machine with good ingredients is called a sound argument. Of course if you don't start with good ingredients, deductive arguments don't guarantee a good end product. Invalid arguments are not 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 lead to a 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 a guarantee. In sum:

INVALID ARGUMENT
Put in false premise(s) → get out either true or false conclusion
Put in true premise(s) → get out either true or false conclusion

VALID ARGUMENT
Put in false premise(s) → get out either true or false conclusion
Put in true premise(s) → get out only true conclusion

Soundness.

To say an argument is valid, then, is not to say that its conclusion must be accepted as true. The conclusion must be accepted *only if* (1) the argument is valid *and* (2) the premises are true. This combination of valid argument plus true premises (and therefore true conclusion) is called a 'sound' argument. Calling it sound is the highest endorsement one can place on an argument. If you accept an argument as sound you are really saying that you must accept its conclusion. This can be shown by the use of another valid, deductive argument. If you say that an argument is sound you are saying two things that may be understood as premises:

1. If the premises of the argument are true, then the conclusion must also be true. (That is to say, you're maintaining that the argument is valid.)
2. The premises of the argument are (in fact) true.

If you regard these two as premises, you can produce a deductive argument that concludes with certainty:

3. Therefore, the conclusion of the argument is true.

For a deductive argument to pass muster, it must be valid. But being valid is 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 the conclusions – as in our cheese and cats example. Surely it is soundness that matters.

Keep in mind, however, 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 simply 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. They satisfy themselves with starting with truth. The problem is that starting with truth doesn't guarantee that one ends with it.

Furthermore in launching criticism, it is important to grasp that understanding validity gives you an additional tool for evaluating another's position. In criticizing another's 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, criticizing and thinking well. It is an 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*
*Patrick J. Hurley, *A Concise Introduction to Logic*, 7th edn (2000)
Fred R. Berger, *Studying Deductive Logic* (1977)

1.5 Invalidity

Given the definition of a valid argument, it may seem obvious what an invalid one looks like. Certainly, it is simple enough to define an invalid argument: it is one 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.

To be armed with an accurate definition, 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. One needs to understand the definition's full import.

Consider this argument:

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

If you're thinking carefully, you'll probably have 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. And if one 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 some 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 is 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 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 a Muslim or Jew. He or she simply may not like pork sausages but frequently enjoy turkey or beef.

So the fact that Ghandi 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 is 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. In this case, the structure is

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

where X is substituted for 'vegetarian', Y for 'person who does not eat pork sausages' and Z for 'Ghandi'. 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 philosophers call a new 'substitution instance' of the argument form.) If we substitute X for 'Cat', Y for 'meat eater' and Z for 'the president of the US', we get:

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

The premises are true but the conclusion clearly false. Therefore this cannot be a valid argument structure. (You can do this with various invalid argument forms. Showing that an argument form is invalid by substituting sentences into the form that give true premises but a false conclusion is what philosophers call showing invalidity by 'counterexample'. See 3.8)

It should be clear therefore that, as with validity, invalidity is not determined by the truth or falsehood of the premises but by the logical relations between them. This reflects a wider, important feature of philosophy. Philosophy is not just about saying things that are true; it is about making true claims that are grounded in good arguments. You may have a particular viewpoint on a philosophical issue, and it may just turn out that you are right. But, in many cases, unless you can show you are right by the use of 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 is the truth.

See also

1.2 Deduction
1.4 Validity and soundness
1.7 Fallacies

Reading

*Patrick J. Hurley, *A Concise Introduction to Logic*, 7th edn (2000)
*Irving M. Copi, *Introduction to Logic*, 10th edn (1998)

1.6 Consistency

Of all the philosophical crimes there are, the one you really don't want to get charged with is inconsistency. Consistency is the cornerstone of rationality. What then, exactly, does consistency mean?

Consistency is a property characterizing two or more statements. If one holds two inconsistent beliefs, then, at root, this means one is asserting both that X is true *and* X is in the same sense and at the same time not true. More broadly, one holds inconsistent beliefs if one belief contradicts another or the beliefs in question together imply contradiction or contrariety.

In short, two or more statements are *consistent* when it is possible for them all to be true at the same time. Two or more statements are *inconsistent* when it is not possible for them all to be true at the same time.

A single sentence, however, can be self-contradictory when it makes an assertion that is necessarily false – often by conjoining two inconsistent sentences.

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 clearly contradicts the first. I am, in effect, saying both that 'all murder is wrong' and 'not all murder is wrong' – a clear inconsistency.

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 (persons). 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 foetuses.'

Defenders of the permissibility of abortion, on the other hand, may retort that properly understood the two claims are not inconsistent. One 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). Or a defender of abortion might change 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 beliefs that intuitively seem perfectly acceptable yet seem to match the definition of inconsistency given. Two examples might be:

It is raining, and it is not raining.
My home is not my home.

In the first case, the inconsistency may be only apparent. What one may really be saying is not that it is raining and not raining, but rather that it is neither properly raining nor not raining, since there is a third possibility – perhaps that it is drizzling, or intermittently raining – and that this other possibility most accurately describes the current situation.

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 inconsistency to be real, the relevant terms being used must retain the same meaning throughout.

This equivocation in the meanings of the words show that we must be careful not to confuse the logical form of an inconsistency – asserting both X and not X – with ordinary language forms that appear to match it but really don't. Many ordinary language assertions that both X and not X are true turn out, when analysed carefully, not to be inconsistencies at all. So be careful before accusing someone of inconsistency.

But, when you do unearth a genuine logical inconsistency, you've accomplished a lot, for it is impossible to defend the inconsistency without rejecting rationality outright! Perhaps there are poetic, religious and philosophical contexts in which this is precisely what people find it proper to do.

Poetic, religious or philosophical inconsistency?

What about the second example we present above – 'My home is not my home.' Suppose that the context in which the sentence is asserted is in the diary of someone living under a horribly violent and dictatorial regime – perhaps by George Orwell's character Winston Smith in *1984* as he writes in his diary. Literally, the sentence is self-contradictory, internally inconsistent. It seems to assert both that 'This is my home' and that 'This is not my home.' But the sentence also seems to carry a certain poetic sense, to convey how absurd the world seems to the speaker, how alienated he or she feels from the world in which he or she exists.

The Danish existentialist philosopher 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. Existentialist philosopher Albert Camus (1913–60) maintained that there is something fundamentally 'absurd' (perhaps inconsistent?) about human existence.

Perhaps, then, 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 is often the case that one can hold a consistent theory that is inconsistent with another, equally consistent theory, the 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 maintained, it may be possible to develop two or more theories that are (1) internally consistent, (2) yet 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.

Take as an example the so-called problem of evil. How do we solve the puzzle that God is supposed to be good but there is awful suffering in the world? A number of theories can be advanced that may solve the puzzle but are inconsistent with one another. One can hold that God does not exist. Or one 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 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 need to appeal to more than consistency if we are to decide between the 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
3.28 Sufficient reason

Reading

José L. Zalabardo, *Introduction to the Theory of Logic* (2000)
Fred R. Berger, *Studying Deductive Logic* (1977)
Pierre M. M. Duhem, *La théorie physique, son object et sa structure* (1906)

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 presents 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.

All fallacies are instances of faulty reasoning. 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.

One of the most common types of inferential error attributable to the form of argument has come to be known as 'affirming the consequent'. It is an extremely easy error to make and can often be difficult to detect. Consider the following example:

1. If Fiona won the lottery last night, she'll be driving a red Ferrari today.
2. Fiona is driving a red Ferrari today.
3. Therefore Fiona won the lottery last night.

Why is this invalid? It is simply this: as with any invalid argument, the truth of the premises does not guarantee the truth of the conclusion. Drawing this conclusion from these premises leaves room for the *possibility* that the conclusion is false, and if any such possibility exists, the conclusion is not guaranteed.

One can see that such a possibility exists in this case by considering that it is *possible* that Fiona is driving a Ferrari today for reasons *other* than her winning the lottery. Fiona may, for example, have just inherited a lot of money. Or she may be borrowing the car, or perhaps she stole it. (Her driving the Ferrari for other reasons does not of course render the first premise false. Even if she's driving it because she in fact inherited a lot of money, it still might be true that *if* she had instead won the lottery she *would have* gone out and bought a Ferrari just the same. Hence the premises and conclusion might all be true but the conclusion will not *follow with necessity* from the premises.

The source of this fallacy's persuasive power lies in an ambiguity in ordinary language concerning the use of 'if'. The word 'if' is sometimes used to imply 'if and only if' but sometimes means simply 'if'. Despite their similarity, these two phrases have very different meanings.

As it turns out, the argument would be valid if the first premise were stated in a slightly different way. Strange as it may seem, while the argument about Fiona above is deductively invalid, substituting either of the following statements for the first premise in that argument will yield a perfectly valid argument.

If Fiona's driving a red Ferrari today, she won the lottery last night.
Only if Fiona won the lottery last night will she be driving a red Ferrari today.

Because 'if' and 'only if' are ordinarily used in rather vague ways (that don't distinguish the usages above), philosophers redefine them in a very precise sense. Mastering philosophical tools will require that you master this precise usage as well (see 4.5).

In addition, because fallacies can be persuasive and are so prevalent, it will be very useful for you to acquaint yourself with the most common fallacies. (The masked man [3.17] and genetic [3.12] fallacies have their own entries in this book. Others are delineated in the texts listed below.) Doing so can inoculate you against being taken in by bad reasoning. It can also save you some money.

Informal fallacies.

The 'gambler's fallacy' is both a dangerously persuasive and a hopelessly flawed species of inference. The fallacy occurs when someone is, for example, taking a bet on the tossing of a fair coin. The coin has landed heads up four times in a row. The gambler therefore concludes that the next time it is tossed, it is more likely to come up tails than heads (or the reverse). But what the gambler fails to realize is that each toss of the coin is unaffected by the tosses that have come before it. No matter what has been tossed beforehand, the odds remain 50-50 for every single new toss. The odds of tossing eight heads in a row are rather low. But if seven heads in a row have already been tossed, the chances of the sequence of eight in a row being completed (or broken) on the next toss is still 50-50.

What makes this an informal rather than a formal fallacy is that we can actually present the reasoning here using a valid form of argument.

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, I'm due for a tails.
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 flaw here is not with the *form* of the argument. The form is valid; logicians call it *modus ponens*, the way of affirmation. It is the same form we used in the valid Fiona argument above. Formally, it looks like this:

If P, then Q.
P
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 all individual tosses) is and remains 50-50 no matter what toss or tosses preceded it. But people mistakenly believe that past flips of coins somehow affect future flips. There's no formal problem with the argument, but because this factual error remains so common and so easy to commit, it has been classified as a fallacy and given a name. It is a fallacy, but only informally speaking.

Sometimes ordinary speech deviates from these usages. Sometimes any widely held, though false, belief is described as a fallacy. Don't worry. As the philosopher Ludwig Wittgenstein said, language is like a large city with lots of different avenues and neighbourhoods. It is all right to adopt different usages when one inhabits different parts of the city. Just keep in mind where you are.

See also

1.5 Invalidity
3.12 Genetic fallacy
3.17 Masked man fallacy
4.5 Conditional/Biconditional

Reading

*S. Morris Engel, *With Good Reason: An Introduction to Informal Fallacies*, 5th edn (1974)
*Irving M. Copi, *Informal Fallacies* (1986)
*Patrick J. Hurley, *A Concise Introduction to Logic*, 7th edn (2000)

1.8 Refutation

Samuel Johnson was not impressed by Bishop Berkeley's argument that matter did 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. Johnson's refutation wildly missed Berkeley's point, since the bishop would never have denied that one could kick a stone. But not only did Johnson's refutation formally fail; it also contained none of the hallmarks of a true refutation.

To refute an argument is to show that is wrong. If one merely disagrees with an argument or denies its soundness, one is not refuting it, although 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. One can show that the argument is invalid: the conclusion does not follow from the premises as claimed (see 1.5). One can show that one or more of the premises are false (see 1.4).

A third way is to show that the conclusion must be false and that therefore, even if one can't identify what is wrong with the argument, something must be wrong with it (see 3.23). This last method, however, isn't strictly speaking a refutation, as one has failed to show *what* is wrong with the argument, only *that* it must be wrong.

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. We may be justified in rejecting an argument even if we have not strictly speaking refuted it. We may not be able to show that a key premise is false, for example, but we may believe that it is 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 no good reasons for believing it to be true and good grounds for supposing it to be false. Therefore we can regard any argument that depends on this premise as dubious and rightly ignore it.

Conceptual problems.

More contentiously, we 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, one might argue that the government is only obliged to provide assistance to those who do not have enough to live on. But given that there can be no precise formulation of what 'enough to live on' is, any argument must be inadequate that concludes by making a sharp distinction between those who have enough and those who don't. The logic of the argument may be impeccable and the premises may appear to be true. But if we use vague concepts in precise arguments we inevitably 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 know the clear difference between refutation and other forms of objection and to be clear what form of objection one is offering.

See also

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

Reading

*Theodore Schick, Jr, and Lewis Vaughn, *How to Think about Weird Things: Critical Thinking for a New Age*, 3rd edn (2002)

1.9 Axioms

Obtaining a guaranteed true conclusion in a deductive argument requires both (1) that the argument be valid, and (2) that the premises be true. Unfortunately, the procedure for determining whether or not a premise is true is much less determinate than the procedure for assessing an argument's validity.

Defining axioms.

Because of this indeterminacy, 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 certain kind of rational system. Axiomatic systems were first formalized by the geometer Euclid (*fl.* 300 BCE) in his famous work the *Elements*. Axioms in such systems are initial claims that stand in need of no justification – at least from within the system. They are simply 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 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 does not 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 this. The problem is that no argument is going to be able to run very far with such an axiom. The axiom is purely tautological, that is to say, 'unmarried men' merely repeats 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.) It is thus a spectacularly uninformative sentence (except to someone who doesn't know what 'bachelor' means') and is therefore 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 parts of mathematics and geometry rest on their axioms, and it is only by accepting their basic axioms that more complex proofs can be constructed. For example, it is an axiom of Euclidean geometry that the shortest distance between any two points is a straight line. But while these axioms are vital in geometry and mathematics, they merely define what is true within the particular system of geometry or mathematics to which they belong. Their truth is guaranteed, but only in the context within which they are defined. Used in this way, their acceptability rises or falls with the acceptability of the theoretical system as a whole. (One might call these propositions, 'primitive' sentences within the system.)

Axioms for all?

Some may find the contextual rendering of axiom we've given rather unsatisfactory. Are there not any 'universal axioms' that are both secure and informative in all contexts, 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 were 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 axiom states that 'if there be no determinate cause it is impossible that an effect should follow' (*Ethics*, bk 1, pt 1, axiom 3). But as John Locke (1632–1704) pointed out, this claim is, taken literally, pretty uninformative since it is 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.

Hume, however, points out that we have no reason to accept this claim about the world. That is to say, we have no reason to believe that events can't occur without causes (*Treatise*, bk 1, pt 3, §14). Certainly, by definition, an effect must have a cause. But for any particular *event*, we have no reason to believe it has followed necessarily from some cause. Medieval Islamic philosopher al-Ghazali (1058–1111) advanced a similar line (*The Incoherence of the Philosophers*, 'On Natural Science', Question One ff.)

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, self-evident truths that may serve as axioms in our reasoning. 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 well. 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

1.1 Arguments, premises and conclusions
1.10 Definitions
1.12 Tautologies, self-contradictions and the law of non-contradiction
6.6 Self-evident truths

Reading

Euclid, *Elements*
Al-Ghazali, *The Incoherence of the Philosophers*
*Benedictus Spinoza, *Ethics* (1677)

1.10 Definitions

If, somewhere, there lies 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 is very easy to argue at cross-purposes or to commit fallacies involving equivocation. As the exploits of a recent US president show, if you are, for example, to debate the ethics of 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 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 ambiguity as possible, and adequate definitions are the perfect tool for helping us do that.

Free trade example.

For example, I may be discussing the justice of 'free trade'. In doing so I may define free trade as 'trade that is not hindered by national or international law'. By doing so I have fixed the definition of free trade for the purposes of my discussion. Others may argue that they have a better, or alternative, definition of free trade. This may lead them to reach different conclusions about its justice. Setting out definitions for difficult concepts and reflecting on their implications comprises a great deal of philosophical work.

The reason why it is 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*. My clear definition of how I will use the term thereby both helps and constrains my discussion. It helps it because it gives a determinate and non-ambiguous meaning to the term. It limits it because it means that what I 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 realizing it, meaning different things by their terms.

Too narrow or too broad?

This is why it is 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 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, 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 criticizing someone's position on some philosophical topic is to indicate a case that fits the definition he or she is using but which they would clearly not wish to include under it.

A definition is like a property line; it establishes the limits marking those instances to which it is proper to apply a term and those instances to which it is not. The ideal definition permits application of the term to just those cases to which it should apply – and to no others.

A rule of thumb.

It is generally 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. However, there will be occasions where it is appropriate, even necessary, to coin special uses. 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 is not enough just to say, 'By knowledge I mean . . .' Rather, the search is for a definition that best articulates the concept in question. Much of the philosophical work along these lines has involved conceptual analysis or the attempt to unpack and clarify the meanings of important concepts. What is to count as the best articulation, however, requires a great deal of debate. Indeed, it is a viable philosophical question as to whether such concepts actually can be defined. For many ancient and medieval thinkers (like Plato and Aquinas), formulating adequate definitions meant giving verbal expression to the very 'essences' of things – essences that exist independently of us.

Many more recent thinkers (like some pragmatists and post-structuralists) have held that definitions are nothing more than conceptual instruments that organize our interactions with each other and the world, but in no way reflect the nature of an independent reality.

Some thinkers have gone so far as to argue that all philosophical puzzles are essentially rooted in a failure to understand how ordinary language functions. While, to be accurate, this involves attending to more than just definitions, it does show just how deep the philosophical preoccupation with getting the language right runs.

See also

1.9 Axioms
3.4 Category mistakes
3.9 Criteria

Reading

*Plato (*c.*428–347 BCE), *Meno*, *Euthyphro*, *Theatetus*, *Symposium*
J. L. Austin, *Sense and Sensibilia* (1962)
Michel Foucault, *The Order of Things* (1966)

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 with absolute certainty. The bedrock was an idea that could not be doubted, the *cogito* ('I think') – *je pense donc je suis* ('I think therefore I am', popularly rendered *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.

Ancient Stoics like Cleanthes (*ob.* 232 BCE) and Chrysippus (280–207 BCE) maintained that we experience certain impressions of the world and morality that we simply cannot doubt – experiences they called 'cataleptic impressions'. Later philosophers like the eighteenth century's Thomas Reid (1710–96) believed that God guarantees the veracity of our cognitive faculties. His contemporary Giambattista Vico (1688–1744) reasoned that we can be certain about things human but not about the non-human world. More recently the Austrian philosopher Ludwig Wittgenstein (1889–1951) tried to show how it simply makes no sense to doubt certain things.

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. 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 described 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 doing so simply renders a psychological account of the concept. It fails to define when we are *warranted* in feeling this way. A more philosophical account of certainty would add the claim that a proposition is certainly true when it is impossible for it to be false – and 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'.

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 skilful 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, most would agree that absolute certainty in advancing truth claims remains unattainable. Moreover, even if achieving this sort of certainty were possible, while it may be that all that's philosophically certain is true, clearly not all that's true is certain.

But if you can't have demonstrable certainty, what is the next best thing? To give a proper answer to this question would require a much larger study of the theory of knowledge. But it is worth saying a little about the answer that most commonly springs to mind: probability.

Probability is the natural place to retreat to if certainty is not attainable. As a refuge, however, it is rather like the house of sticks the pig flees to from his house of straw. The problem is that probability is a precise notion that cannot be assumed to be the next best thing to certainty.

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 radioactive atom, the probability of it having decayed 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 have decayed over the period of the element's half-life and five will not have decayed. On at least some interpretations, it is genuinely indeterminate which atoms will fall into which category.

Subjective probability 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 I 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 is a head and a 50 per cent probability that it is a tail. If you could see the coin, you would know that, in fact, it was 100 per cent certain that it was whichever side was facing up.

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.

Certainty and validity.

If you have a sound deductive argument, then its conclusion follows from the premises with certainty. Many inquirers, however, demand not only that conclusions *follow* with certainty but that the conclusions themselves be 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. All humans are mortal.
2. Socrates was a human.
3. Therefore, Socrates was mortal.

The conclusion of the first argument clearly enters only a probable claim. The conclusion of the second argument also follows with certainty from the premises, but, in contrast to the first, it enters a much more definite 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 is merely probable while the other is not.

You must therefore distinguish between (1) whether or not the conclusion of an argument *follows* from the premises with certainty, and (2) whether or not the conclusion of an argument advances a *statement* whose truth is itself certain.

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?

Some philosophers would say no. For example, they would say that although there remains a great deal of dispute, there is near unanimous agreement among philosophers on many things – for example, that Plato's theory of forms is false and that 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 be seen as a matter of philosophical maturity.

See also

1.1 Arguments, premises and conclusions
1.2 Deduction
1.4 Validity and soundness
1.5 Invalidity
1.9 Axioms

Reading

*Brad Inwood and Lloyd P. Gerson, *Hellenistic Philosophy: Introductory Readings*, 2nd edn (1988)
Giambattista Vico, *Scienza nuova* (1725)
Ludwig Wittgenstein, *On Certainty* (1969)

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

Tautologies and self-contradictions 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 are 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: 'A 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 in logic such that it will turn out to be true in every circumstance – or, as some say, in every possible world. Tautologies are '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, since 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. That is to say, every valid argument may be articulated as a statement of this form: 'If W, X, Y are true, then C is true', where W, X and Y are the argument's premises and C is the 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 is the foundation upon which all logic is built. But, in fact, it is not the law itself that's informative so much as any attempt to break it.

Attempts to break the law of non-contradiction are themselves contradictions, and they are obviously and in all circumstances wrong. A contradiction flouts the law of non-contradiction, since to be caught in a contradiction is to be caught asserting both that something is true and something is false at the same time – asserting both P and not-P. Given that the law of non-contradiction is a tautology, and thus in all circumstances true, there can be nothing more clearly false than something that attempts to break it.

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 any attempt to refute it presupposes it. 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 criticizing 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 a principle that cannot be rationally criticized, because it is 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
3.19 Paradoxes
3.16 Leibniz's law of identity
3.27 Self-defeating arguments

Reading

*Patrick J. Hurley, *A Concise Introduction to Logic*, 7th edn (2000)
Aristotle, *Posterior Analytics*, bk 1, ch. 11:10
Aristotle, *Interpretation*, esp. chs 6–9