JULIAN BAGGINI AND PETER S. FOSL

THE PHILOSOPHER'S
TOOLKIT

A Compendium of Philosophical Concepts and Methods

SECOND EDITION

Julian Baggini is editor and co-founder of *The Philosophers' Magazine* (www.philosophersmag.com). He is the author of several books, including *The Ethics Toolkit* (with Peter S. Fosl, Wiley-Blackwell, 2007), *Welcome to Everytown: A Journey into the English Mind* (2008), *Complaint* (2008) and *Should You Judge This Book by Its Cover?* (2009). He has written for numerous newspapers and magazines, including the *Guardian*, the *Financial Times*, *Prospect* and the *New Statesman*, as well as for the think tanks the Institute of Public Policy Research and Demos.

Peter S. Fosl is Professor of Philosophy at Transylvania University in Lexington, Kentucky. He is co-author with Julian Baggini of *The Ethics Toolkit* (Wiley-Blackwell, 2007) and is also co-editor of the *Dictionary of Literary Biography* (2002) volumes on British philosophy, as well as co-editor with David E. Cooper of *Philosophy: The Classic Readings* (Wiley-Blackwell, 2009). Fosl's scholarly publications address topics in scepticism, ethics, the philosophy of religion and the history of philosophy.

Praise for the first edition

'*The Philosopher's Toolkit* provides a welcome and useful addition to the introductory philosophy books available. It takes the beginner through most of the core conceptual tools and distinctions used by philosophers, explaining them simply and with abundant examples. Newcomers to philosophy will find much in here that will help them to understand the subject.'

David S. Oderberg,
University of Reading

'. . . the average person who is interested in arguments and logic but who doesn't have much background in philosophy would certainly find this book useful, as would anyone teaching a course on arguments, logic, and reasoning. Even introductory courses on philosophy in general might benefit because the book lays out so many of the conceptual "tools" which will prove necessary over students' careers.'

About.com

'Its choice of tools for basic argument . . . is sound, while further tools for argument . . . move through topics and examples concisely and wittily . . . Sources are well chosen and indicated step by step. Sections are cross-referenced (making it better than the Teach Yourself "100 philosophical concepts") and supported by a useful index.'

Reference Reviews

JULIAN BAGGINI AND PETER S. FOSL

THE PHILOSOPHER'S TOOLKIT

A Compendium of Philosophical Concepts and Methods

SECOND EDITION

For Rick O'Neil, colleague and friend, in memoriam

Contents

Alphabetical Table of Contents

Preface

Philosophy can be an extremely technical and complex affair, one whose terminology and procedures are often intimidating to the beginner and demanding even for the professional. Like that of surgery, the art of philosophy requires mastering a body of knowledge, but it also requires acquiring precision and skill with a set of instruments or tools. *The Philosopher's Toolkit* may be thought of as a collection of just such tools. Unlike those of a surgeon or a master woodworker, however, the instruments presented by this text are conceptual – tools that can be used to analyse, manipulate and evaluate philosophical concepts, arguments and theories.

The *Toolkit* can be used in a variety of ways. It can be read cover to cover by those looking for instruction on the essentials of philosophical reflection. It can be used as a course book on basic philosophical method or critical thinking. It can also be used as a reference book to which general readers and more advanced philosophers can turn in order to find quick and clear accounts of the key concepts and methods of philosophy. The aim of the book, in other words, is to act as a conceptual toolbox from which all those from neophytes to master artisans can draw instruments that would otherwise be distributed over a diverse set of texts and require long periods of study to acquire.

For this second edition, we have expanded the book from six to seven sections, and reviewed and revised every single entry. These sections progress from the basic tools of argumentation to sophisticated philosophical concepts and principles. The text passes through instruments for assessing arguments to essential laws, principles and conceptual distinctions. It concludes with a discussion of the limits of philosophical thinking.

Each of the seven sections contains a number of compact entries comprising an explanation of the tool it addresses, examples of the tool in use and guidance about the tool's scope and limits. Each entry is cross-referenced

to other related entries. Suggestions for further reading are included, and those particularly suitable for novices are marked with an asterisk. There is also a list of Internet resources at the back of the book.

Becoming a master sculptor requires more than the ability to pick up and use the tools of the trade: it requires flair, talent, imagination and practice. In the same way, learning how to use these philosophical tools will not turn you into a master of the art of philosophy overnight. What it will do is equip you with many skills and techniques that will help you philosophize better.

Acknowledgements

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 ways 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, 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 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 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 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 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 nonstandard 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 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, literally, that with which an argument concludes, the product and result of an inference or a chain of inferences, 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'. Such unstated premises are called *enthymemes*, and arguments which employ them are *enthymemetic*.

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. 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 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, such as '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 (religious philosophers), or because we grasp them using cognitive faculties certified by God (Descartes, Reid, Plantinga). In our own 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 both to spot the explicit premises and to 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, however, intellectual or otherwise, the construction will only stand if the foundations are secure.

SEE ALSO

READING

★ Nigel Warburton, *Thinking From A to Z*, 2nd edn (2000)
★ Graham Priest, *Logic: A Very Short Introduction* (2001)
    Patrick J. Hurley, *A Concise Introduction to Logic*, 10th edn (2007)

---------------------   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, we confidently 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 is definitely 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 (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 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 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

   Fred R. Berger, *Studying Deductive Logic* (1977)
★ John Shand, *Arguing Well* (2000)
   A. C. Grayling, *An Introduction to Philosophical Logic* (2001)

───────────────── 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 wran-
gling with a dimension of logic philosophers call 'induction'. Unlike deduc-
tive 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 rea-
soning 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 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 problem of induction

Inductive generalizations are, however, often where the action is. Reasoning in experimental science, for example, often 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. Francis Bacon (1561–1626) argued persuasively for just this conception of induction. 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 *definitely* or 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:

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 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. 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 across space and time. 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. (Perhaps to be precise we should only count observed events in the *present*, especially when claims about the past also rely on assumptions about the uniform operations of nature, for example memory.) 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. (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 only *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?

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 accept the uniformity (or at least general 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 generalizations. They are (or at least have been so far!) simply too useful to refuse. Inductive generalizations 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 generalization requires that we hold an indispensable belief which itself, however, must remain in an important way ungrounded.

SEE ALSO

READING

★ Francis Bacon, *Novum Organum* (1620)
★ David Hume, *A Treatise of Human Nature* (1739–40), Bk 1
   Colin Howson, *Hume's Problem: Induction and the Justification of Belief* (2003)