Also by Stephen Law
The Philosophy Files

The Philosophy Gym

25 SHORT ADVENTURES IN THINKING

Stephen Law
ILLUSTRATED BY DANIEL POSTGATE

Thomas Dunne Books
St. Martin's Press New York

THOMAS DUNNE BOOKS.
An imprint of St. Martin's Press.

www.stmartins.com

Library of Congress Cataloging-in-Publication Data

Law, Stephen.
The philosophy gym: 25 short adventures in thinking / Stephen Law; illustrated by Daniel Postgate.—1st. U.S. ed.
p. cm.
Includes index (p. 286).
ISBN 0-312-31452-3
1. Philosophy—Miscellanea. 2. Imaginary conversations. I. Title.

B68.L383 2003
100—dc21 2003050699

First published in Great Britain by REVIEW, an imprint of Headline Book Publishing

First U.S. Edition: December 2003

10 9 8 7 6 5 4 3 2 1

For Tilda

14

WHY EXPECT THE SUN TO RISE TOMORROW?

PHILOSOPHY GYM CATEGORY
WARM-UP
MODERATE
MORE CHALLENGING

Every morning we expect the sun to appear over the horizon. But according to the philosopher David Hume (1711–76), our expectation is wholly irrational. This chapter gets to grips with Hume's extraordinary argument.

An Absurd Claim?

The scene: MacCruiskeen, a scientist, is watching the sunrise. She's accompanied by her close friend Pluck, a student of philosophy.

Pluck: Beautiful sunrise.

MacCruiskeen: Yes. And right on time, too.

Pluck: Yet there was no good reason to expect it to rise this morning.

MacCruiskeen: But the sun has risen every morning for millions of years. Of course it was going to rise this morning as well.

Pluck: There's no reason to suppose it will rise tomorrow, either. In fact, it's just as sensible to expect that a huge million-mile-wide bowl of tulips will appear over the horizon instead.

MacCruiskeen: I agree we can't be *certain* the sun will rise tomorrow. Some cataclysmic event might destroy the earth before then. But it's very *unlikely* that anything like that will happen. The *probability* is that the sun will rise, surely?

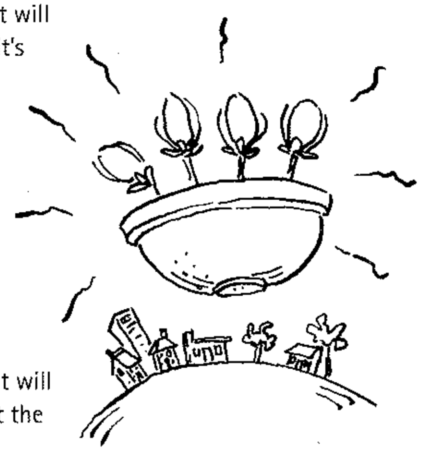

A TULIP SUNRISE

Pluck: You misunderstand me. I'm not just saying we can't be certain the sun will rise tomorrow. I'm saying *we have no more reason to suppose that it will rise than we have to suppose that it won't.*

MacCruiskeen: That's absurd. The evidence – such as the fact that the sun has risen every morning for millions of years – overwhelmingly supports my belief that the sun will rise tomorrow, too.

Pluck: You're mistaken.

Pluck's position might seem ridiculous. But Hume has an argument that appears to show that she's right. Not only is our belief that the sun will rise tomorrow wholly unjustified, but so, too, are all our scientific theories.

Before we look at Hume's argument, I need briefly to explain the difference between deductive and inductive reasoning.

Thinking Tools: Inductive and Deductive Reasoning

An *argument* consists of one or more claims or *premises* and a *conclusion* arranged in such a way that the premises are supposed to *support* the conclusion. Arguments come in one of two forms: *deductive* and *inductive*.

1. Deductive arguments
Here is an example of a deductive argument:

- All cats are mammals.
- My pet is a cat.
- Therefore my pet is a mammal.

Two things are required for a good deductive argument. First of all, the premises must be true. Secondly, the argument must be *valid*. The expression 'valid', in this context, means that the premises must *logically entail* the conclusion. In other words, to assert the premises but to deny the conclusion would be to involve oneself in a *logical contradiction*. The above argument is valid. A person who claims that all cats are mammals and that their pet is a cat but who also denies their pet is a mammal has contradicted him or herself.

2. Inductive arguments
Suppose you observe a thousand swans and discover them all to be white.

You don't come across any non-white swans. Then surely you have pretty good reason to conclude that all swans are white. You might reason like this:

- Swan 1 is white.
- Swan 2 is white.
- Swan 3 is white . . .
- Swan 1,000 is white.
- Therefore all swans are white.

This is an example of an inductive argument. Inductive arguments differ from deductive arguments in that their premises are supposed to *support*, but *not* logically entail, their conclusions. The above argument is not, and is not intended to be, deductively valid. To assert that the first thousand swans examined are white but that not *all* are white is not to contradict oneself (in fact, not all swans *are* white: there are black swans from New Zealand).

Nevertheless, we suppose that the fact that if all the swans we have observed so far are white, then that makes it *more likely* that all swans are white. The premises *support* the conclusion. We believe that an inductive argument can *justify* belief in its conclusion, despite not providing a logical guarantee that if the premises are true then the conclusion will be.

Why Is Induction Important?

We rely on inductive reasoning in arriving at beliefs about what we have not observed, including, most obviously, our beliefs about what will happen in the future.

Take, for example, my belief that the next time I sit in a chair it will support my weight. How is this belief justified? Well, I have sat in a great many chairs and they have always supported my weight before. That leads me to think it likely that the next chair I sit in will support my weight, too.

But notice that the statement that all the chairs I have ever sat in have supported my weight does not *logically entail* that the next chair will. There is no *contradiction* in supposing that even though I have never before experienced a chair collapse beneath me, that is what's about to happen.

But it then follows that I can't justify my belief that the next chair will not collapse by means of a *deductive* argument from what I have observed. So *if my belief is justified at all, it must be by means of an inductive argument.*

Science is heavily dependent on induction. Scientific theories are supposed to hold *for all times and places*, including those we have not observed. Again, the only evidence we have for their truth is what we have observed. So, again, we must rely on inductive reasoning to justify them.

The Unjustified Assumption

We have seen that inductive reasoning is important. Science depends on it. If it can be shown that inductive reasoning is wholly irrational, that would be a catastrophic result. Yet that's precisely what Hume believes he can show.

Let's return to Hume's argument. Hume believes it is no more rational to suppose the sun will rise tomorrow than it is to suppose that it won't. Hume's argument, in essence, is simple: it's that *induction rests on a wholly unjustified and unjustifiable assumption.* What is this assumption? Pluck proceeds to explain.

Pluck: Your belief that the sun will rise tomorrow is irrational. Hume explained why. Whenever you reason to a conclusion about what you haven't observed, you make an *assumption*.

MacCruiskeen: What assumption?

Pluck: You assume that *nature is uniform*.

MacCruiskeen: What do you mean?

Pluck: I mean you assume that those patterns that we have observed locally are likely to carry on into those portions of the universe that we haven't observed, including the future and the distant past.

MacCruiskeen: Why do I assume that?

Pluck: Well, put it this way: if you *didn't* believe that nature is uniform, then the fact that the sun has, in your experience, risen every day *wouldn't* lead you to expect it to continue to rise, would it?

MacCruiskeen: I guess not.

Pluck: So you see – *it's only because you assume nature is uniform that you conclude that the sun will continue to rise in the future.*

It appears that Pluck is right. Whenever we reason inductively, we make an assumption about the uniformity of nature. We assume that the universe is patterned throughout in just the same way.

Imagine an ant sitting in the middle of a bedspread. The ant can see that its bit of the bedspread is paisley-patterned. So the ant assumes the rest of the bedspread – the bits it can't see – are paisley-patterned, too. But why assume this? The bedspread could just as easily be a patchwork quilt. The bedspread could be paisley here, but plaid over there and polka-dotted over there. Or perhaps, just over the ant's horizon, the print on the bedspread turns to a chaotic mess, with blobs, lines and spots muddled up quite randomly.

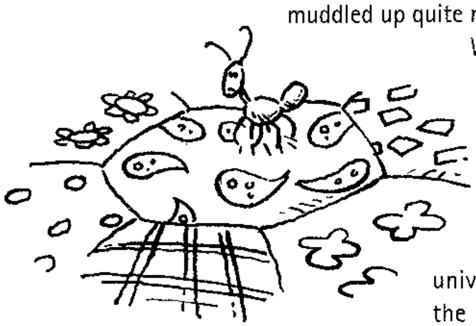

We are in a similar position to the ant. The universe could also be a huge patchwork, with local regularities, such as the ones we have observed – the sun rising every day, trees growing leaves in the spring, objects falling when released, and so on – but no *overall* regularity. Perhaps the universe becomes a chaotic mess just over the horizon, with events happening entirely randomly. What reason have we to suppose this isn't the case?

As Pluck is about to explain, it seems we have none.

Pluck: So the problem is this: unless you can *justify* your assumption that nature is uniform, your use of induction is itself unjustified. But then so, too, are all those conclusions based on inductive reasoning, including your belief that the sun will rise tomorrow.

MacCruiskeen: True.

Pluck: So how do we *justify* the assumption that nature is uniform?

We have just two options: we can either appeal to *experience* – to what you have observed – or you might try to justify the assumption *independently of experience*. MacCruiskeen is happy to admit that we cannot know that nature is uniform without observing nature.

MacCruiskeen: Obviously, we can't know independently of experience that nature is uniform.

Pluck: I agree. Our five senses – sight, touch, taste, hearing and smell – provide our only window on the world. Our knowledge of nature is dependent on their use.

MacCruiskeen: True.

Pluck: Which means that, if the assumption that nature is uniform is to be justified at all, it must be by appeal to what we have *experienced* of the world around us.

MacCruiskeen: Yes. But *isn't* the claim that nature is uniform justified by experience?

Pluck: No. To say that nature is uniform is to make a claim about what holds for *all* times and places.

MacCruiskeen: True.

Pluck: But you can't directly observe *all* of nature, can you? You can't observe the future. And you can't observe the distant past.

MacCruiskeen: Also true.

Pluck: But then your justification of the claim that nature is uniform must take the following form. You observe nature is uniform *around here* at the *present time*. Then you *infer* that nature is also like that at all those other times and places. Correct?

MacCruiskeen: I suppose so.

Pluck: But that is *itself* an inductive argument!

MacCruiskeen: Yes, it is.

Pluck: Your justification is, therefore, *circular*.

Here we reach the nub of Hume's argument. It seems that, if it can be confirmed at all, the assumption that nature is uniform can only be confirmed by observing that nature is uniform *around here* and then concluding that this is what it must be like *overall*.

But such a justification would itself be inductive. We would be using precisely the form of reasoning we're supposed to be justifying. Isn't there something unacceptably circular about such a justification?

The Circularity Problem

Pluck certainly thinks so.

MacCruiskeen: What is the problem with the justification being circular?

Pluck: Look, imagine that I think The Great Mystica, the psychic who works at the end of the pier, is a reliable source of information.

MacCruiskeen: That would be very foolish of you!

Pluck: But suppose my justification for trusting The Great Mystica is that she claims to be a reliable source of information. I trust her because she says she's trustworthy.

MacCruiskeen: That would be no justification at all! You need some reason to suppose that The Great Mystica is trustworthy *before* you trust her claim that she is.

Pluck: Exactly. Such a justification would be unacceptably circular because it would *presuppose* that The Great Mystica was reliable.

MacCruiskeen: I agree.

Pluck: But your attempt to justify induction is unacceptable for the very same reason. To justify induction you must first justify the claim that nature is uniform. But in attempting to justify the claim that nature is uniform you rely on induction. That won't do. You're just *presupposing* that induction is reliable.

We can now sum up Hume's extraordinary argument. All inductive reasoning, it seems, relies on the assumption that nature is uniform. How, then, might this assumption be justified? Only by experience, surely. But we cannot *directly observe* that nature is uniform. So we must *infer* that it is uniform from what we have directly observed: that is, from a *local* uniformity. But *such an inference would itself be inductive*. Therefore we cannot justify the assumption. So our trust in induction is unjustified.

'But Induction *Works*, Doesn't It?'

Perhaps you're not convinced. You might suggest that there is one very obvious difference between, say, trusting induction and trusting The Great Mystica. For induction *actually works*, doesn't it? It has produced countless true conclusions in the past. It has allowed us successfully to build supercomputers, nuclear power-stations and even to put a man on the moon. The Great Mystica, on the other hand, may well have a very poor track record of making predictions. That's why we are justified in believing that induction is a reliable mechanism for producing true beliefs, whereas trusting The Great Mystica is not.

The problem, of course, is that this is itself an example of inductive reasoning. We are arguing, in effect, that induction has *worked until now*, and therefore induction will continue to work. Since the reliability of induction is what is in question here, it seems that this justification is, again, unacceptably circular. It is, after all, just like trying to justify trust in the claims of The Great Mystica by pointing out that she herself claims to be reliable.

An Astonishing Conclusion

The conclusion to which we have been driven is a sceptical one. Sceptics claim that we do not know what we might think we know. In this case the scepticism concerns *knowledge of the unobserved*. Hume and Pluck seem to have shown that we have no justification for our beliefs about the unobserved, and thus no *knowledge* of the unobserved.

Hume's conclusion is a fantastic one. It's a good test of whether someone has actually understood Hume's argument that they acknowledge its conclusion is fantastic (many students new to philosophy misinterpret Hume: they think his conclusion is merely that we cannot be *certain* what will happen tomorrow). In fact, so fantastic is Hume's conclusion that MacCruiskeen cannot believe that Pluck is really prepared to accept it.

MacCruiskeen: You're suggesting that what we've observed to happen so far gives us *no clue at all* as to what will happen in the future?

Pluck: Yes. Things *may* continue in the same manner. The sun may continue to rise. Chairs may continue to support our weight. But we have *no justification whatsoever* for believing any of these things.

MacCruiskeen: Let me get this straight. If someone were to believe that it's just as likely that a huge bunch of tulips will appear over the horizon tomorrow morning, that chairs will vanish when sat on, that in future water will be poisonous and objects will fall upwards when released, we would ordinarily think them *insane*. Correct?

Pluck: Yes, we would.

MacCruiskeen: But if you're right, these 'insane' beliefs about the future are actually just as well supported by the available evidence as is our 'sensible' belief that the sun will rise tomorrow. Rationally, we should accept that these 'insane' beliefs are actually just as likely to be true!

Pluck: That's correct.
MacCruiskeen: You *really* believe that? You really believe it's just as likely that a million-mile-wide bowl of tulips will appear over the horizon tomorrow morning?
Pluck: Well, actually, no, I don't.
MacCruiskeen: Oh?
Pluck: I do believe the sun will rise tomorrow. For some reason, I *just can't help myself*. I see that, *rationally*, I *shouldn't* believe. But while I realise that my belief is wholly irrational, I can't stop believing.

Hume's Explanation of Why We Believe

Like Pluck, Hume admitted that we *can't help* but believe that the sun will rise tomorrow, that chairs will continue to support our weight, and so on. In Hume's view, our minds are so constituted that when we are exposed to a regularity, we have no choice but to believe the regularity will continue. Belief is a sort of involuntary, knee-jerk response to the patterns we have experienced.

Thinking Tools: Reasons and Causes – Two Ways of Explaining *Why* People Believe What They Do

Hume's explanation of why we believe that the sun will rise tomorrow does not, of course, give us the slightest reason to suppose that this belief is actually *true*.

It is useful to distinguish two very different ways in which we can 'give the reason' why someone believes something. We may give the *grounds* or *evidence* that a person has for holding a belief. Or we may explain what has *caused* this person to believe what they do.

It's important to realise that *to offer a causal explanation of a belief is not necessarily to offer any sort of rational justification for holding it*. Consider these explanations:

- Tom believes he is a teapot because he was hypnotised during a stage act.
- Anne believes in fairies because she is mentally ill.
- Geoff believes in alien abduction because he was indoctrinated by the Blue Meanie cult.

These are purely causal explanations. To point out that someone believes they are a teapot because they were hypnotised into having that belief during the course of a hypnotist's routine is not to provide the slightest grounds for supposing that this belief is true.

The following explanation, on the other hand, gives the subject's grounds for belief (which is not yet to say they are good grounds):

- Tom believes in astrology because he finds that newspaper astrology predictions are quite often correct.

Interestingly, ask the hypnotised person why they believe they are a teapot and chances are they will be unable to answer. The correct *causal* explanation is unavailable to them (assuming they don't know they have been hypnotised). But nor will they be able to offer a convincing *justification* for their belief. They may simply find themselves 'stuck' with a belief that they may themselves recognise is irrational.

Hume admits that, similarly, his explanation of why we believe the sun will rise tomorrow does not supply the slightest grounds for supposing that this belief is true. Indeed, we have no such grounds. It is, again, a belief we simply find ourselves 'stuck' with.

Conclusion

If Hume is right, the belief that the sun will rise tomorrow is as unjustified as the belief that a million-mile-wide bowl of tulips will appear over the horizon instead. We suppose the second belief is insane. But if Hume is correct, the first belief is actually no more rational. This conclusion strikes us as absurd, of course. But Hume even explains *why* it strikes us as absurd: we are made in such a way that we *can't help* but reason inductively. We can't help having these irrational beliefs.

Hume's argument continues to perplex both philosophers and scientists. There's still no consensus about whether Hume is right. Some believe that we have no choice but to embrace Hume's sceptical conclusion about the unobserved. Others believe that the conclusion is clearly ridiculous. But then the onus is on these defenders of 'common sense' to show precisely *what* is wrong with Hume's argument. No one has yet succeeded in doing this (or at least no one has succeeded in convincing a majority of philosophers that they have done so).

What to read next

This chapter contains an argument for the existence of God – the 'argument from miracles'. Other arguments for God's existence can be found in Chapter 7, Does God Exist?, and Chapter 1, Where Did the Universe Come From?.

Further reading

An excellent discussion of miracles, the supernatural and all things weird is provided by:

Theodore Schick Jr and Lewis Vaughn, *How to Think about Weird Things*, second edition (California: Mayfield, 1999).

Simon Blackburn provides a succinct introduction to Hume's thinking on miracles in:

Simon Blackburn, *Think* (Oxford: Oxford University Press, 1999), Chapter 5.

24

HOW TO SPOT EIGHT EVERYDAY REASONING ERRORS

PHILOSOPHY GYM CATEGORY
WARM-UP
MODERATE
MORE CHALLENGING ■

A fallacy is an error in reasoning. Reason – the use of argument – is the main tool of the philosopher. But, of course, we also depend on reason in our everyday lives. So it's important that we can spot a logical howler when we come across it.

This chapter will help you to identify eight common reasoning errors (errors that, very probably, you sometimes make, too).

1. The Post Hoc Fallacy (a Fallacy of the Superstitious)

I had been worried about my exams. So Jill bought me a rabbit's foot to take with me for luck. I took the foot, and I passed the first exam. So, you see, the rabbit's foot worked! I shall take it to all my other exams, and it will make me pass them, too.

This is an example of the post hoc fallacy. Here are two more examples:

- John's psychic told him she would send positive psychic vibes when he tried to climb Everest. And he succeeded! So, you see, his psychic really does have miraculous powers! From now on he's always going to ask her for help in climbing mountains.
- Local taxes went up. And, look, the crime figures went up. So higher local taxes cause crime. Local taxes should never have been raised!

Examine all three examples and you will find that someone concludes that, because one event occurred *after* another, therefore the first event must have *caused* the second.

This is clearly flawed reasoning. Usually, when one event occurs after another, there is no causal connection between them. Suppose, for example, that I plug in the kettle. Immediately after, a comet crashes into Jupiter. Did I cause the comet's impact? Obviously not.

Of course, there *may* be a causal connection between two events that occur one after the other. Perhaps the rise in taxes really did cause a rise in crime. Perhaps John's psychic really did cause him to succeed. The point is that such 'one-off' observations do not remotely *justify* the claim that the first event caused the second.

The moral is: *don't leap to conclusions*. Noticing that one event occurs immediately after another might give one grounds for investigating whether two events are causally connected. But it does not, by itself, make it rational to believe that there is any such connection.

Unfortunately, superstitious people are very prone to the post hoc fallacy, and the unscrupulous can and do take advantage. Point out that just after someone bought one of your lucky rabbit's feet they immediately won some money on a scratch card and you will soon find gullible customers beating a path to your rabbit's foot store.

2. Argument from Authority (a Favourite of Celebrity Advertisers)

- 'I'm going to find my perfect partner soon.' 'How do you know?' 'I consulted the fortune-telling machine on the pier, and it said so.'
- 'Blancmange face packs are an effective beauty treatment.' 'How do you know?' 'All the celebrities are using them – Anita Sopwith Camel, actress and pop star, even advertises them on TV.'
- 'Genetic engineering is always morally wrong; it should never be carried out.' 'Why do you believe that?' 'Because Dr Bits told me.' 'Is Dr Bits an expert in ethics and genetic techniques?' 'No, he's a professor of mathematics.'
- 'I believe that Brand X washes whiter than any other brand.' 'Why?' 'Because scientists working for the Brand X corporation say so.'

Sometimes we're justified in believing something because an authority on the subject tells us that it is true. If a professor of chemistry warns you not to drop a lump of phosphorus into a sink full of water, I would follow her advice.

But often such 'appeals to authority' are fallacious.

In the first two examples, the 'authorities' in question are highly dubious. Why should a celebrity be any better informed about the efficacy of blancmange face packs than anyone else?

In the third, while Dr Bits really is an authority, he is not an authority on the issue in question. There is no reason to suppose that his opinion on the ethics of genetic engineering is any more reliable than anyone else's.

In the fourth example, the authority in question may be biased. To what extent can we trust scientists working for a particular company to give impartial advice about its products?

When appealing to a supposed 'authority', you must be warranted in supposing that it really is an authority on the issue in question, that there aren't many other authorities on the issue holding an opposing view, that the authority is not significantly biased, and so on. *Only then* is it sensible to place your trust in the authority in question.

3. Slippery Slope (the Miser's Favourite)

- If I lend you one pound today, tomorrow it will be two pounds, then ten pounds. Pretty soon you will owe me thousands!

This is an example of the slippery slope fallacy. It occurs when someone argues that one thing will inevitably follow from another but without providing any justification for supposing that 'slide' from one thing to the other is likely to happen. Usually, there are a number of intermediate steps involved in the 'slide'.

Is the following an example of the fallacy?

- If we allow someone to select the sex of their baby today, tomorrow we will have to allow selection for eye and hair colour. Pretty soon, we will have to allow 'designer babies'.

Yes, it is, if no reason is given for supposing that we cannot or will not simply stop at some point along the 'slide'.

4. False Dilemma (the Salesperson's Favourite)

It is common to argue like this:

- Either *A* or *B*. Not *A*. Therefore *B*.

This is often a perfectly acceptable form of argument, as in this case:

- Either John has a driving licence or else John is not permitted to drive. John has not got a driving licence. Therefore John is not permitted to drive.

This argument, on the other hand, is not acceptable:

- Either 1 + 1 = 5 or 2 + 2 = 5. It is not true that 1 + 1 = 5. Therefore 2 + 2 = 5.

Why not? Because, unlike in the first argument, the alternatives presented in the either/or premise could *both* be false. People often construct such arguments without registering that there might be other alternatives, as in this example:

- Either we cut welfare or the government goes into the red. We cannot allow the government to go into the red. Therefore we must cut welfare.

In this case, there are other options not mentioned, such as raising taxes. Customers are often railroaded into making bad decisions by a salesperson's use of false dilemma:

- Either you give a substantial donation to the Blue Meanie cult or you will have an unhappy life. You don't want an unhappy life, do you? So make that donation!
- Either you buy the Kawazuki K1000 for great home sound entertainment, or else you make do with second-rate rubbish. Are you really prepared to accept second-rate rubbish? I thought not. So you have no choice, do you? You have to buy the Kawazuki K1000!

Be cautious when salespeople appear to offer you an inescapable either/or decision. As often as not they are using false dilemma.

5. Trying Only to Confirm (a Favourite of Politicians the World Over)

Suppose I show you four cards, each of which has a letter on one side and a number on the other. An 'E', 'F', '2' and '5' are visible, like this:

Now suppose I ask you what is the quickest way of establishing that the following is true: *of the four cards shown, those with vowels on one side have even numbers on the other.* Which cards do you need to turn over to establish that the hypothesis is true? Take a moment to think about it . . . Probably you think the E and 2 cards should be flipped. In fact, that's actually the wrong combination of cards to turn over. Yet *most people* believe that E and 2 are the cards to examine (so did I when I first saw this test).

So what cards should you flip? The answer is E and 5. Why?

You need to turn the E card over to check that there is an even number on the reverse. If there isn't, the hypothesis is false. You also need to turn the 5 card to check that it doesn't have a vowel on the reverse. If it does, the hypothesis is false. As long as E has an even number and 5 doesn't have a vowel, the hypothesis is true. It doesn't matter what's on the reverse of the F and the 2.

So why are we led astray? Why do we tend to turn the 2 and not the 5? It seems *we have an in-built tendency to try to confirm such hypotheses rather than disconfirm them.* We turn the 2 because we are searching only for positive instances of the hypothesis, not negative ones. We tend to look for confirming evidence, even when a search for disconfirming evidence might be far more telling. This tendency can lead us into serious trouble. Here's another example.

A politician believes that cutting local taxes will cause the crime rate to drop. So she asks her researchers to look for examples of situations where local taxes were cut and the crime rate fell. They find that there are a hundred such examples. So the politician concludes that she is justified in supposing that by lowering local taxes she can cut crime.

The politician sought only to confirm her hypothesis, not disconfirm it. That may have led her astray. Had her researchers bothered to look, they might have found *two* hundred cases in which the crime rate went up after local taxes were cut.

The moral is: when testing a hypothesis, make sure you look not just for confirming evidence, but also for disconfirming evidence.

6. The Gambler's Fallacy

Here are two examples of the gambler's fallacy.

Simon: Still buying those scratch cards?
Stan: Yes. I've been playing for three years and I haven't won yet.
Simon: So why do you bother?
Stan: Well, as I haven't won yet, I must be due for a win fairly shortly!

Tracey: Did you win anything at the dogs last night?
Bob: No. I bet on Rover Dover three times in a row and he lost each time.

Tracey: So you won't bother betting on him again, I guess?
Bob: I shall definitely bet on him again! You see, his record shows that he wins fifty per cent of his races. As he has lost the last three, it follows that he *must* win the next three to even things up! Rover Dover is now a dead cert!

In each case, someone takes the probability of an event A happening over a period of time, notices that, over the first part of that period, the actual incidence of A is much lower than what is probable, and then concludes that A must be much more probable over the rest of the period. They predict a short-term increase in the probability of A to 'even things up' over the longer term.

The fallacy can also work the other way: someone might suppose that a higher-than-expected incidence of A must result in a short-term lowering in the probability of A to 'even things up', as in this case.

Ruth: Doing the lottery again this week?
John: Yes. What numbers are you going to pick?
Ruth: H'm. Well, the numbers that have come up most are 3, 7 and 28. So I certainly shan't be choosing them. As they have come up a lot recently, they are bound not to come up again for a quite a while.

The gambler's fallacy is extremely common. Wait a few minutes around any lottery or scratch card outlet and it won't be long before you hear someone saying that they are 'due' a win, that they won't make the mistake of picking the same numbers that won last week, and so on.

The truth, of course, is that it makes not one jot of difference what has happened up to now. Each week, the probability of any particular sequence of numbers coming up in the UK lottery is always exactly the same: about 14 million to one.

Interestingly, I recently saw a news reporter commit just this fallacy on TV. A couple who chose the same numbers week after week in the UK lottery forgot to buy a ticket the very week that those numbers came up. The couple were devastated, but insisted they would keep on choosing the same numbers in future. The reporter concluded that, sadly, the couple were now far less likely to win with those numbers.

7. Circular Justification (also Known as 'Begging the Question')

Tom: The Great Mystica is a reliable source of information.
Sarah: How do you know?
Tom: She told me so herself.

Bert: God must exist.
Ernie: Why?
Bert: It says so in the Bible.
Ernie: How do you know the Bible is reliable?
Bert: Because it is the word of God.

Violet: John is honest.
William: How do you know?
Violet: Tom told me.
William: How do you know Tom is honest?
Violet: Jane told me.
William: How do you know Jane is honest?
Violet: John told me.

Each of these justifications runs in a circle. In each case, the truth of the claim that is supposed to be justified is actually *assumed* by that justification. Such circular justifications are unacceptable: you can't justify a claim simply by assuming it to be true.

8. The Fallacy of Affirming the Consequent

Take a look at the following argument:

- If I am a man, then I am mortal. I am a man. Therefore I am mortal.

There's nothing wrong with this argument. It has two premises, both of which are true. And the conclusion follows. Now look at these arguments:

- If John is happy, then John is playing football. John is playing football. Therefore John is happy.

- If I am taller than Sue, then Sue is short. Sue is short. Therefore I am taller than Sue.

Are these arguments acceptable? Interestingly, a study of people who had no training in logic found that *over two-thirds believed arguments of this form to be acceptable*. Yet both arguments are faulty. Each resembles the first argument we looked at, but differs from it in an important way. The first argument has this form:

- If A, then B. A. Therefore B.

The faulty arguments have this form:

- If A, then B. B. Therefore A.

It is known as the fallacy of *affirming the consequent*. To work through a concrete example, look again at my first illustration of the fallacy above. It is true that if John is happy, then John is playing football. Football is the only thing that makes John happy. Does it follow that if John is playing football, then he is happy? No. For while John may be happy only when playing football, it may also be that he is often unhappy even when he *is* playing football.

Here, finally, are a couple of philosophical examples of affirming the consequent:

- If God exists, then there is good in the world. There is good in the world. Therefore God exists.
- If other humans feel pain, then they will cry out when injured. Other humans will cry out when injured. Therefore other humans feel pain.

Further reading

This chapter provides just a few examples of fallacies. For more examples, see:

Nigel Warburton, *Thinking from A to Z* (London: Routledge, 1996).

There is a useful list of fallacies with both explanation and examples at:

www.nizkor.org/features/fallacies/

25

SEVEN PARADOXES

PHILOSOPHY GYM CATEGORY
WARM-UP
MODERATE
MORE CHALLENGING

This chapter contains seven of the most famous, fascinating and infuriating paradoxes. All the examples in this chapter take the form of seemingly plausible arguments leading to seemingly implausible conclusions. They leave us flummoxed because, while we are unwilling to accept the conclusion, we can't see anything wrong with the reasoning that leads us to the conclusion.

See if you can figure out solutions to the following seven examples. But be warned: some of the world's greatest minds have tried and failed. Indeed, the first of our paradoxes is alleged to have caused the early death of Philetas of Kos.

Many readers will be content to dip into my seven examples just for fun: they are curiously entertaining. Others may wish to pursue things further. For the second group, I have included some further hints and comments at the end.

Paradox 1: The Man who Spoke the Truth but Didn't

A traveller was walking one day when he met an old man sitting beside the road smoking a pipe.

'The first thing said to you by the first person you meet today will not be true,' said the old man. 'Trust me – don't believe what he says!'

'OK,' said the traveller. 'But hang on a minute: *you're* the first person I've met today.'

'Exactly!' said the old man.

You may have spotted something funny going on here. If the old man speaks the truth, then the first thing he says is not true. But if the first thing he says is not true, then the first thing he says is true.

This is a version of the famous *liar paradox*, a paradox first formulated in ancient Greece over 2,000 years ago.

The traveller thought he saw a way out of the paradox: claim that what the old man first said is *neither true nor not true*. After all, why *does* every such sentence have to be either true or not true?

'Old man, you're trying to trick me,' said the traveller. 'It's obvious that what you said is *neither* true nor not true.'

'Aha,' said the old man. 'You're suggesting that it is not true that what I said is true, and also not true that what I said is not true?'

'That's exactly right,' said the traveller.

'Well, then, if it's not true that what I said is true, then what I said *is* not true!'

The traveller was starting to get a headache. The old man continued: 'And if it's not true that what I said is not true, then what I said *is* true! For what I said is precisely that what I said is not true!'

The traveller was starting to feel like ramming the old man's pipe down his throat.

'So you see,' said the old man, 'your suggestion is wrong: it's *not* true that what I said is neither true nor not true. In fact, it's *both* true and not true!'

But that's impossible. Isn't it?

Paradox 2: The Sorites Paradox

Here are two versions of this ancient paradox.

Jenny's Sandpit

Jenny is tidying her sandpit while Jim looks on.

'You know, the ants from that ants' nest over there keep stealing grains of your sand.'

Jenny looked down at the line of ants. Each marched up to her heap, took a single grain of sand between its mandibles and carried it off down the garden.

Jenny didn't seem much bothered.

'But they'll never be able to remove this heap of sand, will they?' she replied.

'Why not? Look, if they keep on removing those grains one by one, then eventually there will be just a single grain left, won't there? It might take weeks, but eventually you'll have just a single grain of sand left at the bottom of your pit. Then you won't have a heap of sand any more, will you?'

Jenny scratched her head. 'But look, *by removing a single grain of sand from a heap, you can't turn it into a non-heap, can you?*'

'No, obviously not,' replied Jim. 'For example, if I have 1,000 grains, and I remove one grain, giving me 999 grains, I still have a heap. Correct?'

'Right,' said Jenny. But then, no matter how many grains are removed by those ants, they will *never* succeed in turning my heap into a non-heap.'

Jim was now very confused. 'But if that's true, then a single grain of sand *is* a heap!'

'Precisely!' said Jenny. 'In fact, even *no* grains of sand is a heap!'

But it's surely false that no grains of sand is a heap. So where did Jenny go wrong?

Bob's Balding Spot

Bob was looking forlornly into the bathroom mirror while holding a pocket mirror up to the back of his head.

'There goes another hair,' he said sadly.

'Stop worrying,' replied Sarah. 'You can't turn from being not bald to being bald with the loss of a single hair, can you?'

'I guess not,' said Bob.

'So you're still not bald, are you?' said Sarah.

'I suppose not. But hang on! If what you say is true, then, no matter how many hairs fall out of my head, I will never be bald!'

'Er. I didn't say that.'

'But it does *follow* from what you said, doesn't it? Suppose there are exactly a million hairs on my head now, and I'm not bald. If one hair is removed, and you're right that removing a single hair can't transform a non-bald person into a bald one, then I still won't be bald. Remove another hair, and I still won't be bald. Remove yet another, and I still won't be bald. And so on, until there are no hairs left. I still won't be bald! But clearly I *will* be bald! So it follows that your principle that, by removing a single hair from his head, you can't turn a person from being not bald into being bald, must be false!'

'You're mad.'

'But it follows! In fact, there must come a point where, by losing just a *single hair*, I'll turn from being bald into being not bald!'

'But that's absurd. There's not a precise number of hairs that marks the boundary between being bald and being not bald.'

'But there must be!'

'But then *what is* that number of hairs?'

'I don't know. Maybe it's 10,027. Maybe it's 799. But there must be such a number.'

'That's just plain silly.'

'Actually, it *must* be true! In fact, perhaps the hair that just fell out was the one that turned me from being not bald into being bald!'

Paradox 3: The Boastful Barber

Luigi, the barber of Seville, was proudly boasting of his success. 'You know, I'm the man who shaves all and only those men in Seville who don't shave themselves!'

'I can't believe that,' says Franco.

'Why not?'

'Well, do you shave yourself? If you do, then, from what you just said, it follows that you *don't* shave yourself. For you said you shave all *and only* those who don't shave themselves. Right?'

'Right. But what if I tell you that I don't shave myself – my wife does the job for me?'

'Well, if you don't shave yourself, then it follows that you do. For you said you shave *all* and only those who don't shave themselves. Right?'

'Er. Right.'

So does Luigi just shave those who don't shave themselves? Or doesn't he?

Paradox 4: Achilles and Tortoise

Achilles rides a huge motorbike. Tortoise has a little moped. They decide to have a race. But as his motorbike is much faster than Tortoise's moped, Achilles decides to give Tortoise a head start.

Achilles starts at A. Tortoise starts at B. By the time Achilles has made up the distance to B, Tortoise has moved forward to C. By the time Achilles reaches C,

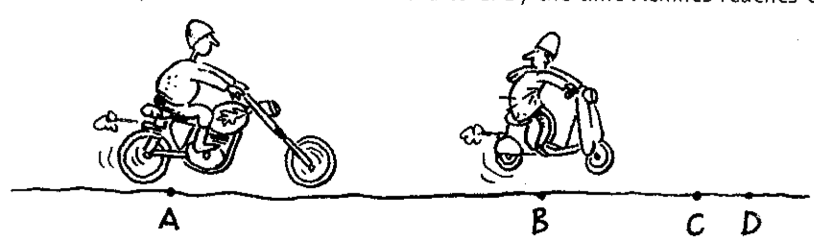

Tortoise has got to D. Every time Achilles manages to close the gap between where he is and where Tortoise is, Tortoise has moved forward a bit more. But there are going to be an infinite number of such gaps to close before Achilles finally catches up with Tortoise. But one can never travel across an *infinite* number of gaps, for no matter how many gaps one travels across there will always be an infinite number yet to travel. There's no last gap. *So Achilles can never catch Tortoise.*

Yet, of course, he can. How come?

Paradox 5: The Ravens

Pluck is asking Bridie, a scientist, what it is that scientists do.

Pluck: How does science work?

Bridie: Well, scientists construct theories that are confirmed by observation.

Pluck: Give me an example.

Bridie: Very well. Take the generalisation that all ravens are black. Now *all generalisations are confirmed by their instances.* So, for example, an observation of a black raven, being an instance of the generalisation that all ravens are black, confirms that generalisation. Each black raven one sees confirms the hypothesis that all ravens are black a little more.

Pluck: I see. But look, it's true, is it not, that if two hypotheses are logically equivalent, then whatever confirms one hypothesis should confirm the other?

Bridie: That must be true. Logically equivalent hypotheses are really just two different ways of saying the same thing. So whatever confirms one hypothesis should confirm the other.

Pluck: Right. But the hypothesis that *all ravens are black* is logically equivalent to the hypothesis that *all non-black things are non-ravens.*

Bridie: True. In effect, they say the same thing.

Pluck: But then if all generalisations are confirmed by their instances, then a non-black non-raven confirms that all non-black things are non-ravens, right?

Bridie: True.

Pluck: But then a non-black non-raven confirms that all ravens are black, true?

Bridie: Er. True, I suppose.

Pluck: So white shoes, red poppies and blue skies – being non-black non-ravens – all confirm the hypothesis that all ravens are black.

Bridie: But that's absurd!

Pluck: But *it follows from what you agreed before.* My observation of this pink blancmange confirms that all ravens are black!

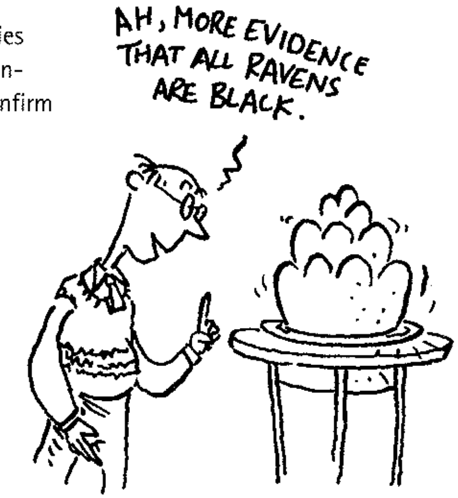

Pluck is right: if everything Bridie agreed to is correct, then a pink blancmange really does confirm that all ravens are black. But that's absurd. Or is it?

Paradox 6: The Unexpected Examination

The teacher tells her students that they should expect an exam some time during the next week. But she does not say when. The exam is to be unexpected.

But will it be unexpected?

Can the teacher give the exam on Friday? No, for if she gives it on Friday, then on Friday morning her students, knowing that they hadn't had the exam earlier in the week, will expect it.

What about Thursday? Well, as her students know that the exam cannot be on Friday, if she leaves it till Thursday, the exam will again be expected. So Thursday is out, too.

What about Wednesday? That, again, is ruled out. Her students know the exam cannot be on either Thursday or Friday, so if the teacher leaves it until Wednesday, the exam will again be expected.

But then Tuesday and Monday are also out, and for the same reason.

In short, *the teacher cannot give an unexpected examination.*

Yet, of course, she can. Can't she?

Paradox 7: 'Santa Claus Doesn't Exist'

Little Brian is reading an English grammar book.

'Dad. Names are used as labels for people and other things, aren't they?'

'That's right. The job of a name within a sentence is always to pick out someone or something so that you can then go on to say something about it.'

'Right. So if I say "Jack is tall", what I say is true when the person the name "Jack" refers to has the property of being tall, and false otherwise.'

'You've got it.'

'But wait a second. Yesterday you said "Santa Claus doesn't exist", right?'

'I did, yes.'

'And what you said is true?'

'Of course.'

'But how *can* it be true? You said the job of a name in a sentence is to pick out an object so that one can then go on to say something about it. But the name "Santa Claus" doesn't pick out anyone, does it?'

'Er, no.'

'But then "Santa Claus" cannot do its job within the sentence, can it? In which case, the sentence "Santa Claus doesn't exist" cannot be true, can it?'

'H'm. I guess not.'

'But you just said it *is* true.'

Little Brian's question is a good one. How can 'Santa Claus doesn't exist' be true if the job of the name 'Santa Claus' within a sentence is to refer?

General Advice for Solving Paradoxes

Here's a hint on how to solve paradoxes. All the paradoxes in this chapter take the form of *arguments*. An argument is made up of one or more claims or *premises* and a *conclusion*. The premises are supposed to *support* the conclusion.

These arguments are paradoxical because the premises are plausible, the conclusion implausible, yet the reasoning apparently cogent.

When faced with such a paradox, you always have three options:

- Explain why at least one of the premises of the argument seems true but is false.
- Explain why the conclusion of the argument seems false but is true.
- Expose some logical flaw in the reasoning.

Before you do any of these things, however, it often helps to identify and set out the argument clearly. This is not always as easy as it sounds.

Here, by way of illustration, is the sandpit version of the sorites paradox set out in more formal style (I assume, for the sake of argument, that Jenny's sandpit contains a heap of 100,000 grains).

- If *n* number of grains is a heap, then so is *n*–1 grains. 100,000 grains of sand is a heap. Therefore, 99,999 grains is also a heap.

This same form of reasoning is then reapplied over and over again (dropping the figures in the middle premise and conclusion by one each time) until you reach the conclusion that zero grains of sand is a heap.

Your options are: 1. to accept the conclusion, 2. to reject the reasoning, or 3. to reject one of the premises.

Here are some further tips and comments on the seven paradoxes we have looked at.

Paradox 1

There's no consensus on how this paradox should be solved. You might, perhaps, be tempted just to bite the bullet and say: 'OK, so what the old man says is both true and not true. It's a contradiction. What's the problem with admitting the existence of contradictions?'

This strategy won't work. Not only are there plenty of problems with admitting contradictions (which I won't go into here), but we can in any case rework the paradox so that admitting contradictions doesn't help.

Here's how. Suppose we introduce the prefix 'UN-*P*' in such a way that 'UN-*P*' applies to all and only those things to which the term '*P*' applies. That's just stipulated. So, for example, 'UN-horse' applies to all and only those things that aren't horses. Now consider this sentence:

This sentence is UN-true.

It follows that this sentence is both true and UN-true. But we just defined 'UN-' in such a way that, *by stipulation*, nothing can be both true and UN-true. Admitting contradictions does nothing to solve *this* version of the paradox.

Paradox 2

Again, there's no agreement about how to solve this paradox. Some philosophers insist that there must be a precise number of grains of sands marking the boundary between a heap and a non-heap. So it's not true that removing a single grain will never turn a heap into a non-heap. It's just that we don't know what this precise number is.

But the suggestion that there is such a precise boundary is a lot to swallow. Surely, it is *we* who determine what our concepts are and where their boundaries lie. So how could our concept of a heap come to have a precise boundary of which we are ignorant?

Paradox 3

This paradox is fairly easy to solve: deny that there is any such person as Luigi, the barber who shaves all and only those who don't shave themselves. So the sentence 'Luigi shaves just those who don't shave themselves' is neither true nor false.

Paradox 4

Here's a similar paradox.

Movement is impossible. For suppose I wish to move one yard. In order to move one yard, I must move half that distance: half a yard. But to move half a yard I must move a quarter-yard, and so on ad infinitum. So I must make an infinite number of movements before I can move a yard. But I cannot make an infinite number of movements, for an infinite number of movements can never be completed. Therefore, I cannot move one yard (or even a bit of a yard).

Paradox 5

One of the more popular strategies here is to deny the principle that all generalisations are confirmed by their instances. And in fact there are other counterexamples to this principle. Take the generalisation that *all snakes are located outside Ireland*. An instance of this would be: *Fred is a snake and Fred is located outside Ireland*. But the more such instances one accumulates – the more snakes outside Ireland one observes – surely the more likely it is that there *are* snakes in Ireland. So our generalisation about snakes is actually *disconfirmed* by its instances!

Paradox 6

You should assume two things for this paradox to work: that the students can be pretty sure there will be an exam (otherwise even on Friday the exam might be unexpected: they might believe the teacher to have forgotten all about it, and when she doesn't forget that might come as a surprise), and that the students are rational and have good memories (they won't simply forget about the exam, or get confused, so that it does come as a surprise).

Paradox 7

This paradox continues to perplex philosophers of language. Note that it won't do to say that the name 'Santa Claus' *doesn't* refer to a *person* but *does* refer to *something*: it refers to our *concept* of Santa Claus. The reason this won't do is that it would then follow that, as our *concept* of Santa Claus *does* exist, so 'Santa Claus does not exist' would be false.

Further reading

For a clear, rigorous and entertaining introduction to paradoxes, I recommend:

Michael Clark, *Paradoxes from A to Z* (London: Routledge, 2002).