

# *Preface*

I once overheard a telling conversation between two of my colleagues. One asked the other about a new book on a topic of some importance to both of them. He asked whether they would have to do anything different because of the book. The second colleague said not, so the first colleague said he would not read the book. The conversation encapsulates an excellent test of the worth of a philosophical work: an idea is important if as a result of it experts will have to change what they do when they work on the idea's topic. It's not good enough to be right, or to choose an important topic, or to choose a topic of contemporary interest, or to be original; it isn't even enough to have all of those qualities. I know now that much of the work I have done, especially as a graduate student, failed that test—even work that found its way into leading philosophy journals.

However, due in part to some dumb luck I have stumbled across what I believe to be a worthy idea, one that epistemologists cannot ignore. It is, in effect, a new kind of argument for a new kind of radical scepticism. It must be admitted that the notion of scepticism elicits strange behaviour in philosophers, especially epistemologists. Many philosophers, even contemporary ones who should know better, sometimes assert that no one is really a sceptic. Philosophers are pretty much professionally *forbidden* from being radical sceptics even though we aren't forbidden from believing any of many other comparably outlandish claims (e.g., see chapter 5 and part B of chapter 11 of this essay). Scepticism is commonly, and justifiably, thought to be an epistemic pit with particularly nasty properties. If you are in that pit, so is everyone else; there is no practical way of coming out of the pit and into the light of knowledge; and you most likely have a serious cognitive or epistemic deficit, some kind of failure of your epistemic systems. But, as I will presently show, while these claims hold for traditional forms of scepticism, none of them apply to the sceptical theories I explore here. The sceptical snares these new theories lay forth apply to just some people (e.g., many philosophers and many scientists), are caused by some pretty random sociological events, are fallen into only in adulthood, can be escaped on one's own or with the help of others, and indicate no epistemic deficit of any significance. In fact, falling into the sceptical snare is actually an improvement in epistemic standing compared to most who have the knowledge denied to those in the snare. That's right: the person who knows is epistemically inferior to the person who doesn't know. If I'm not mistaken, we have an entirely new kind of scepticism.

In this essay I have restricted myself to investigating the new sceptical argument. I have discussed as little as possible traditional arguments for traditional forms of scepticism. I have not presented my best guesses as to how particular anti-sceptics will react to my new sceptical arguments. I do this not out of any disrespect but in order to make the essay relatively short, unified, and not cumbersome. I have ignored most traditional epistemological disputes and theories. I have resisted,

with effort, the temptation to blather on about my own pet peeves and ideas. I have, with the exception of this preface and a few other places, attempted to avoid the word 'I'. I have done all this in the hopes that a short book restricted to presenting the case for the new scepticism is less distracting than a long one that drifts from the main task. It's also easier to write.

This essay would have been an article of reasonable length if I had not been so anti-sceptical by inclination and training. More than half of the essay is devoted to filling out the details of the sceptical solution to the sceptical puzzles generated by my arguments in the first part of the essay (with the intention of finding problems with that sceptical solution), and investigating anti-sceptical responses to those puzzles. The new sceptical arguments conclude that I don't know some very ordinary facts that one would think are very easily known. For instance, I don't know that my shirt is red, that I currently have a throbbing pain in my knee, or even that *I believe* that my shirt is red or that my knee is throbbing. These are difficult results to swallow, certainly quite contrary to common sense, and in this essay I have done my best to avoid swallowing them. As far as I have determined, I have failed. My current best guess is that many philosophers and most members of certain classes of scientists lack these and analogous ordinary pieces of knowledge, if not today then perhaps in our future and definitely in some very realistic possible worlds. Much of the interest in the new kind of scepticism lies in its scope and strength. Regarding strength: *these results remain even if we adopt all the clever anti-sceptical fixes thought up in recent years*: for instance, reliability, proper functioning, relevant alternatives, contextualism, and the rejection of epistemic closure. That is, even if one or more of those anti-sceptical strategies succeeds in defeating the traditional arguments for radical scepticism, the new sceptical arguments go through. Regarding scope: even if my sceptical arguments are sound, you are free to know millions of exotic truths such as the fact that black holes exist! You can know about the existence of black holes, but not about the colour of your shirt or even about what you believe regarding the colour of your shirt. What kind of scepticism is this?

Enough teasing! Here is the brutally short and crude version of my long and exceedingly sophisticated argument. In order to know P one must be able to rule out some  $\neg P$  possibilities. For instance, in order to know that the tree is a fir one has to rule out the possibility that it's a spruce or hemlock. At least, one has to rule out those possibilities provided they are "real, live" scientifically respectable hypotheses, one is aware that they have such respect, and one is perfectly aware that those hypotheses conflict with one's belief that the tree is a fir. For instance, you came across the tree while taking a stroll through a forest with a tree expert. You said the tree was a fir but she said that it's quite hard to tell from this vantage point because spruces and hemlocks look the same and there are lots of them around here. Those  $\neg P$  possibilities are "relevant alternatives", as it is often said. Perhaps the brain-in-a-vat possibility doesn't need to be ruled out; but the spruce and hemlock possibilities do need to be ruled out. Assuming I can't rule out the spruce possibility, I don't know that the tree is a fir—even if the tree is a fir. At least, I don't know it's a fir tree once I'm *aware* of the live status of the spruce and hemlock possibilities.

But now here's the kicker: there are several "real, live" scientifically respectable hypotheses that can be used in the very same argument template as in the previous paragraph. It's a real, live scientifically expert-endorsed possibility—at least in some possible worlds next door to ours—that no ordinary objects are coloured; that no one believes anything; that no one has any character traits; and that all pains are in brains only. Using the line of reasoning from the tree story, we can show that those of us aware of the live status of these four error theories don't know that fire engines are red, don't know that we sometimes have pains in our lower backs, don't know that John Rawls was kind, and don't even know that we *believe* any of those truths. However, people unfamiliar with philosophy and cognitive science do know all those things, as they are well outside the domain of philosophy and cognitive science. And all of this holds on the mere assumption that those error theories are *false* but live. Finally, there is an intriguing way to transform the live sceptic's argument into one for universal scepticism, the thesis that we know absolutely nothing.

"Yes, but those cases are different from the tree case, and those differences will block the route to the new scepticism." Maybe so, but it turns out to be awfully hard to back up this intuition. I have been unable to find any way to justify it—and I'm anti-sceptical by nature. Furthermore, even if the intuition is right we can derive some very interesting epistemological results from the subsequent denial of the new scepticism. So no matter what our reaction to the live sceptic's argument, we walk away with interesting results. In that respect, we all win.

Thanks are due to the University of Leeds for granting me a semester of research leave to work on a boring project that I abandoned in favour of writing this book. Thanks also to a *Noûs* referee (a presentation of the book's main line of argument appears in my "When a Skeptical Hypothesis Is Live", which is forthcoming in *Noûs*), Paul Bloomfield, David Chalmers, Richard Fumerton, Sanford Goldberg, John Greco, Ken Himma, Sarah McGrath, Joseph Melia, Peter Millican, Mark Nelson, Duncan Pritchard, and Tom Stoneham for written or oral comments, and the faculties at the University of Liverpool and University of Connecticut for helpful discussion. Special thanks are due to my generous colleague Andrew McGonigal, with whom I discussed the issues raised in this book on many very fruitful occasions, and my wife Margaret Frances for offering substantive comments on the entire essay. Without their good sense and perceptive comments the book would have had many more mistakes and fewer insights.

*BRSF*

## *Introduction: When a Sceptical Hypothesis Is Live*

Philosophy is inseparable from scepticism, which follows it like a shadow that it chases away by refuting it, only to find it once again under its feet.

Emmanuel Levinas

*Otherwise than Being or Beyond Essence*

When Jo was a teenager she learned that a huge meteorite wiped out the dinosaurs. She learned this theory in the usual way, hearing it from her parents, teachers, and books. Now pretend that at the time she was told the meteor story as a child, say at the age of eight, the scientific community was sharply divided on the issue of what caused the demise of the dinosaurs. Although most scientists accepted the meteor hypothesis, many others subscribed to the idea that their death was caused by some enormous solar flare. A significant number of other scientists thought that it wasn't a solar flare or a meteor but a particularly nasty series of supervolcanos. These latter two classes of dissenters had decent evidence: evidence concerning the sun and supervolcanos that the meteor advocates took seriously. Both the solar flare theorists and the supervolcano theorists were highly respected professors, highly respected by the meteor theorists and at the top of their profession. Whole book series, conferences, and PhD dissertations were devoted to these competing hypotheses. Suppose further that upon going to her university Jo found out about the rival and highly respected hypotheses. She didn't understand all the reasons why they were so well respected and endorsed, but she was well aware that they were well respected and frequently endorsed by the experts, even the best among them. Even so, she kept her meteor belief.

As it turned out, the meteor hypothesis was the right one. But although some experts may have known that fact, surely Jo was in no position to know it after she found out about the eminent status of the rival hypotheses. She could do little or nothing to defeat the rival hypotheses. Even after studying the issues as an undergraduate she couldn't *know* that a meteor wiped out the dinosaurs, for from the moment she first learned of the dinosaurs until she graduated with a Bachelor's degree in palaeontology she was perfectly aware that there were two "live" contrary hypotheses that she was in

no position to rule out—and they had not been ruled out for her, either, for instance by her teachers. And if she was aware of all this and the competing hypotheses weren't ruled out, then she couldn't know that the meteor hypothesis was correct. Perhaps in order to know the meteor truth she wouldn't have needed to rule out the possibility that she is a brain in a vat. And perhaps someone mostly outside the scientific community and thus not exposed to the *ultimately* misleading supervolcano and solar flare evidence could know the meteor hypothesis is true. Still, in order for Jo to know the meteor fact she *does* have to be able to rule out, to some degree anyway, the solar flare and supervolcano possibilities, for, unlike the crazy philosophical hypotheses, these are “real, live possibilities” *and* she is perfectly aware of their existence, live status, and inconsistency with her belief. Perhaps she doesn't have to completely *demolish* the supervolcano and solar flare hypotheses in order to know the meteor story, but she certainly has to knock them down a few epistemic notches.

The heart of the dinosaur argument is simple: because the supervolcano and solar flare hypotheses are real, live possibilities inconsistent with the meteor hypothesis, Jo is *aware* of all that, she is nothing even approaching an expert or genius on these matters, and those hypotheses are not ruled out, she doesn't know that the meteor story is true. There aren't many ideas in philosophy that can't be coherently and cleverly challenged, but it is a brave philosopher who thinks that the above story is misguided. Surely, one would think, in *those* circumstances Jo's true and partially justified belief in the meteor hypothesis didn't amount to knowledge.

The traditional sceptic uses an argument of roughly the same form.

- a. Here's a sceptical hypothesis (e.g., dreaming, evil demon, hallucination, brain in a vat, wildly defective reasoning skills, universe created five minutes ago).
- b. You must be able to rule out that sceptical hypothesis in order to know much of anything.
- c. However, you can't rule it out.
- d. Thus, you don't know much of anything.

This “ruling-out” type of argument is not, of course, the only way to use sceptical hypotheses to generate powerful sceptical arguments, and one can generate such arguments without sceptical hypotheses. But such an argument is not without an intuitive basis.

The typical reaction to this familiar type of argument for scepticism is to first zero in on the conclusion, pronounce it crazy, and then go back to the rest of the argument with the expectation of finding faults that cannot be repaired. This impulse seems retrospectively justified because when we act on it we immediately become highly suspicious of the sceptical hypotheses employed in the arguments. The suspicions of both philosophers and ordinary folk typically stem from three gut reactions to the hypotheses.

- i. I can't do anything to prove or refute the classic sceptical hypotheses *head on*, so to speak, no matter how hard I try. For instance, there's nothing I can *do*, such as pinching myself or performing clever experiments, to defeat or establish the brain-in-a-vat (BIV) hypothesis. So aren't they idling in some important way?
- ii. The sceptical hypotheses just don't seem relevant, for some reason, to our ordinary claims to know. Sure, I guess they're logically inconsistent with what we claim to know, but so what? Why does that fact about mere logical possibility make the sceptical hypotheses threats to my knowledge that my socks are blue?
- iii. No one has even the slightest positive reason to think the sceptical hypotheses are actually true. In fact, for some of them we lack reason to think that they're even *possibly* true. So how could they constitute a hurdle for my knowledge, given there's nothing positive to be said on their behalf?

None of these thoughts constitute objections, as they are just hints at objections, but some or all of them often justifiably incline us to reject the whole lot of hypothesis-based sceptical arguments, no matter how diverse its membership. They justify, to a significant extent, our conveying the following challenge to the sceptic: until you can use sceptical hypotheses that get around these three problems, the sceptical arguments utilizing these weird hypotheses are never going to be good enough to convince us of their highly counterintuitive conclusions. In truth, the analogy between the dinosaur and traditional sceptical arguments is quite weak due to the fact that (i)-(iii) apply to the traditional sceptical story but not to the dinosaur story. That's a large part of why we don't buy the sceptic's reasoning.

In fact, one might think that the lesson of (i)-(iii) is that the traditional sceptical hypotheses never really posed *any* epistemic threat at all to our knowledge, despite the logical inconsistency of the hypotheses with much of our alleged knowledge. On the other hand, one might concede the presence of an epistemic threat generated by the sceptical hypotheses but conclude that because the hypotheses are so crazy or empty or farfetched it is incredibly *easy* to defuse their threat, without even putting any thought into it. If you're sane and breathing, well, then surprisingly enough that's sufficient to defuse the meagre threat posed by the outrageous evil demon hypothesis. These two responses to the sceptic are variations on the fundamental anti-sceptical idea that *traditional sceptical hypotheses don't pose epistemic threats to our beliefs*—either because we have all somehow automatically done something to *render* them non-threatening or because they were never really a threat from the start. So the traditional sceptic almost never converts us, even though she does win our admiration (because we have such a hard time figuring out where she has gone wrong) and gratitude (because she makes us probe the nature of knowledge and epistemic warrant much more deeply than we would without her). This must be frustrating for the sceptic: the permanent loser whom everyone respects but very few accept.

If a sceptic could present sceptical hypotheses that elude (i)-(iii) because they are real, live hypotheses akin to the supervolcano and solar flare theories, her argument would be much stronger than any of the classic sceptical arguments: critics could no longer responsibly start off their investigation with the presupposition that the sceptic is wrong. We would have to abandon our usual, comfortable reactions to the familiar arguments for scepticism, as the new sceptic's arguments would be so crucially different from the traditional ones. And this is exactly what my new sceptic does:

- The sceptical hypotheses she plugs into her sceptical argument template to generate individual sceptical arguments and conclusions are empirically testable and immune to irrelevancy objections. Further, there are compelling scientific and philosophical reasons to think the hypotheses are *actually true*. Each is a “real, live possibility” actually endorsed by some of today's experts: some of them currently think the hypotheses are *true*. As a consequence, for the new sceptical hypotheses it is much harder to reject (b) or (c) in the sceptical argument given above. The sceptical hypotheses are currently live, so most of us can't rule them out. And because they are live and are widely appreciated to be inconsistent with some enormous and cherished set of beliefs, we have to be able to rule them out in order to have those beliefs amount to knowledge. In effect, with the new sceptical hypotheses each of (i)-(iii) above is false, thereby cutting off the initial anti-sceptical reactions suggested above.
- In some respects the “live” sceptic is more radical than most traditional sceptics. She admits that for all she can show, you may know all sorts of exotic things, such as the fact that the universe is expanding; so in that respect her scepticism is limited (but see final point below). However, included in her scepticism are denials of knowledge claims that traditional sceptical arguments usually do not touch (e.g., you often know what you believe).
- The live sceptic announces that although her arguments apply to you, so that many of your seemingly most secure true beliefs fail to amount to knowledge, two friends of yours might actually escape the arguments, and so know millions of the things you don't. If one friend is unacquainted with any philosophy or much science, then perhaps he can keep his knowledge. If the other is a genius, then she can not only retain her extensive knowledge but also know the sad fact that you are a know-nothing loser when it comes to the beliefs covered by the live sceptical arguments. Furthermore, with some work, or appropriate help from others, one can climb out of the sceptical trap and come to know what one could not know before; so one's time in the live sceptical snare might be brief.
- If that were not enough to pique your interest, the new sceptic holds that falling into the sceptical trap, and thus *losing* knowledge, is actually an *improvement in one's epistemic*

*position* and does not reflect any defect in one's cognitive systems, either permanent or temporary.

- Traditional sceptical conclusions have supporting arguments of course, but such arguments have never generated much in the way of a satisfactory explanation of *why* their scepticism is supposed to be true. The live sceptic has a nice, familiar, and detailed explanation of why her sceptical conclusion holds.
- The implications of the new scepticism are circumscribed but surprising: the new sceptical hypotheses don't generate the mere denial of *certain* knowledge or "high standards" knowledge but the entirely modest amount of epistemic warrant we ordinarily expect. In fact, *these new sceptical results remain even if we adopt all the clever anti-sceptical fixes thought up in recent years in the defence of ordinary knowledge claims: reliability, proper functioning, relevant alternatives, contextualism, the rejection of epistemic closure, et cetera* (more on all of these later). That is, even if one or more of those anti-sceptical strategies succeeds in defeating the traditional forms of scepticism, the new sceptical arguments go through.
- Finally, some epistemologists think that if you are fully aware that Q is a necessary condition on anyone's knowing P, and you are fully aware of the fact that there's a real, live possibility that Q is false (so Q isn't something silly like 'You're a brain in a vat'), then in order for you to know P you must be able to rule out the possibility that Q is false. If that's right, then a new sceptical argument finishes with a conclusion truly *universal*: you fail to know P for *any* P, not just the restricted P's—even though your friends still know loads of truths, you are in a better epistemic position than most of them are, and you can climb out of this universal sceptical trap too. As with the other live sceptical arguments, this one employs a sceptical hypothesis that some experts today actually take to be true, utterly unlike the classic sceptical hypotheses.

In this essay I will defend the outrageous thesis that there is such a set of sceptical hypotheses, each of which generates a new sceptical puzzle. Since I will be generating several new sceptical arguments from the same argument template and defending the seven bulleted claims above, you might expect the outline of my essay to be something like this:

1. Present intuitive considerations supporting the general form of the live sceptical arguments, especially regarding the main principles it uses.
2. Articulate the first live sceptical hypothesis and defend the application of the general argument to it.
3. Formulate the other live sceptical hypotheses and sketch how to apply the general argument to them.

4. Explain why the sceptical solution to the live sceptical puzzles—the position that contains the central claim that the sceptical arguments are sound—has the features described in the first five bulleted points above.
5. Show how to construct a new argument for universal scepticism which is loosely based on the general argument template.
6. Examine the kind of sceptical conclusion common to the live sceptical arguments.
7. Examine general anti-sceptical strategies.
8. Articulate and evaluate particular anti-sceptical solutions to the live sceptical puzzles.

You would be right, even regarding the order. But before we get to the live sceptical hypotheses we must first meet the new theory of Smith and Jones.