

Scepticism and Reliable Belief, by José L. Zalabardo. Oxford: Oxford University Press, 2012. Pp. xii + 216. Price £40.00.

Scepticism plays a curious role in epistemology. On the one hand, it is regarded as a disease. It is, at best, false and, at worst, incoherent. On the other, much ink has been spilled in the attempt to defeat scepticism. While scepticism is regarded as best avoided, it isn't that easy to do so. Thus, it is usually seen as a constraint on an acceptable account of knowledge that it avoid scepticism, even at all costs. The problem of scepticism has endured because it isn't clear whether extant accounts of knowledge meet this constraint.

Against this background, José Zalabardo's new book is refreshing. In it he develops a sophisticated reliabilist account of knowledge which borrows from Robert Nozick's classic *Philosophical Explanations* (Harvard University Press, 1981) and Sherrilyn Roush's recent *Tracking Truth* (Oxford University Press, 2007). While reliabilism is introduced as an antidote to scepticism (Chapter 1), Chapters 2-6 argue on independent grounds that his version of reliabilism is the best account of knowledge available. The grounds offered are familiar: it matches (most of) our intuitions, and appeals to notions that can be elegantly cashed out. In Chapter 7 Zalabardo shows that the account doesn't fully deal with the sceptical problem. While our first-order beliefs are safe from sceptical attack, our epistemic assessments of these beliefs aren't - my beliefs often count as knowledge, but my beliefs that those beliefs are true almost always don't. Given the usual constraint on an acceptable account of knowledge, most would conclude that the account has to be rejected. But Zalabardo argues that it is our metaphysical picture that needs to be rejected, not his epistemological picture (Chapter 8). While I would have liked to hear more about these metaphysical issues, it is unavoidable that a book focusing on epistemology had to give them less time.

In the rest of this review I'll outline the shape of Zalabardo's account and discuss the three types of knowledge he identifies (inferential, non-inferential and default). I finish by returning to scepticism.

Zalabardo distinguishes between three types of knowledge. First, we have inferential knowledge. Ignoring certain complexities, a true belief counts as inferential knowledge *iff* it is adequately supported by the evidence. Evidence E adequately supports belief H *iff* the probability of H conditional on E is sufficiently high and the probability of having E conditional on H is significantly higher than the probability of having E conditional on $\sim H$ (pp. 87-9). Second, we have non-inferential knowledge. My true belief H counts as non-inferential knowledge *iff* it meets two conditions: the probability of H conditional on my believing H must be sufficiently high and the probability of my believing H conditional on H must be significantly higher than the probability of my believing H conditional on $\neg H$ (pp. 111-119). The first condition is the probabilistic version of a safety condition, whereas the second condition gives much the same results as a probabilistic version of a sensitivity condition. Third, we have default knowledge, which covers a special sort of belief which counts as knowledge *iff* it is true (pp. 136-9). Examples include my beliefs that sense perception is reliable and that I'm not a brain in a vat.

Zalabardo's view is broadly foundationalist. A piece of default or non-inferential knowledge supports a piece of inferential knowledge *iff* it provides adequate support for it. Because of the conditions he puts on the adequate support relation, Zalabardo denies the transmission principle, which says that, if S knows E, and knows that E logically entails H, then S is in a position to know H on the basis of E (pp. 99-100). (I might infer that I'm not a handless brain in a vat from my belief that I have hands, but my belief that I have hands doesn't provide adequate support for my belief that I'm not a handless brain in a vat). The transmission principle must be

distinguished from the closure principle, which says that, if S knows E, and knows that E logically entails H, then S is in a position to know H. While Zalabardo denies that the closure principle holds in certain cases (pp. 162-165), he thinks it holds in many of the cases in which Nozick denied it, in particular it is not the case that I know I have hands but not that I'm not a handless brain in a vat (this belief counts as default knowledge).

Inferential Knowledge

Zalabardo has to complicate his account of inferential knowledge to deal with familiar problems with Gettier cases, Moorean inferences and bootstrapping. Zalabardo deals with these problems by putting further conditions on inferential knowledge (Chapter 5). The history of attempts to deal with the Gettier problem might make us sceptical about the likelihood of success, but there is a further worry. Zalabardo is engaging in the 'analytical project' of providing necessary and sufficient conditions for a true belief to count as knowledge. A standard worry about this project is that the outcome will inevitably be a set of conditions that are unwieldy, ad hoc and uninformative. The conditions are likely to be unwieldy because of the number of cases that need to be taken into account. They are likely to be ad hoc because the motivation for the conditions is usually dealing with a counter-example. They are likely to be hard to cash out because the notions that usually figure in them (such as 'reliability') are imprecise.

At first glance, Zalabardo faces all of these problems. But things are more complicated. First, each of his conditions are cashed out using notions (such as providing adequate support) that are made precise via probability theory. Chapters 3-6 can be read as an advertisement for the work probability theory can do in formulating an informative version of reliabilism. This isn't a new idea, and Roush's book pursues the same project. But getting the details right is tricky, and it is no objection to the analytical project that it takes work.

Second, while Zalabardo proceeds as if his aim is to fit his account of knowledge around our intuitions, he is open to discarding particular intuitions if they don't fit with his otherwise well-motivated account (see p. 40 and pp. 48-9). This suggests that Zalabardo sees the methodology here as a sort of reflective equilibrium. We start by formulating conditions on knowing that deal with central cases, and then check to see what our conditions say about other, more marginal, cases. If the conditions seem to get the wrong results, we can either refine our conditions or argue that we need to revise our intuitions. While neither point will move someone who is already convinced that the analytical project can't succeed, one wonders whether the idea that the project must fail is much more than an article of faith.

Non-Inferential Knowledge

In according as much importance to sensitivity as to safety Zalabardo's account is at odds with much of the recent literature. Put in probabilistic terms, a sensitive belief is (roughly) a belief that I am very likely to have given it is true, whereas a safe belief is (equally roughly) a belief that is very likely to be true given I believe it. While the probabilistic version of safety faces problems with beliefs that are extremely likely to be true but are intuitively not known (e.g. my belief that my lottery ticket won't win), sensitivity has its own problems. Consider this case, from Jonathan Vogel. It's very hot and a few hours earlier I left an ice-cube outside. Even though I haven't gone to check, intuitively I know that the ice-cube has melted. But my belief that it has melted isn't sensitive, as I am no less likely to believe the ice-cube has melted if it hasn't than I am if it has.

Zalabardo has a couple of moves here, but I will focus on his argument that this is a case of inferential knowledge. The idea is that I form my belief that the ice-cube will melt (MELT) on the basis of a law to the effect that ice-cubes melt at a certain temperature (LAW), and LAW adequately supports MELT. The probability of MELT given LAW is incredibly high and,

because laws are refuted by their counter-instances, the probability of LAW conditional on \neg MELT is zero. While this deals with Vogel's case, one might wonder whether it causes problems with lottery propositions. Consider a standard setup in which I form the belief that my ticket won't win (LOSER) on the basis of my understanding of the probabilities, knowledge that the lottery is fair, etc. (SETUP). But, as Zalabardo argues (pp. 129-33), the crucial difference is that, while the probability of LOSER given SETUP is incredibly high, the probability of SETUP given LOSER is equal to the probability of SETUP given \neg LOSER. The crucial difference between the two cases is that if the ice-cube doesn't melt LAW is false, whereas my ticket winning doesn't change the fact that it was incredibly unlikely to win.

Now consider another case, from Ernie Sosa. I have dropped my rubbish bag down the chute. While I can't see that it has hit the bottom, I have done this many times before, and each time it has hit the bottom. Further, I can see no obstacle that could have prevented the bag from hitting the bottom. Call the propositions that constitute my evidence GEN. Intuitively, so long as the bag hit the bottom, I know that it hit the bottom (call this proposition HIT). Zalabardo's account will secure this result if one of the propositions constituting GEN is an inductive law that would be refuted if HIT were false. But it is hard to see why this would be the case. While the various propositions that constitute GEN might include laws such as the law of gravity (bags that are dropped will fall until something gets in their way), these laws won't be refuted if HIT were false (all that will have happened is that something got in the way). This problem will generalise to all cases of inductive knowledge where the knowledge isn't based on an inductive law that would be refuted if the known proposition were false.

Default Knowledge

Beliefs that count as default knowledge don't track the truth (if I were a handless brain in a vat, I would still believe that I wasn't) and it is likely that we lack adequate evidence for them (p. 138). So why think of default knowledge as a sort of knowledge? While the thought that there is a sort of knowledge not based on evidence is often defended on Wittgensteinian grounds, Zalabardo suggests an evolutionary analogy (p. 137). Herring gulls are very reliable when identifying their own chicks, but very unreliable when identifying their own eggs. The standard evolutionary explanation for this is that, because Herring gulls often find chicks that aren't their own in their nests, they need an ability to identify their chicks, but because they rarely find eggs that aren't their own in their nests they don't need an ability to identify their eggs. Similarly, we have the ability to track the truth for some beliefs (e.g. my belief that I'm writing this review in a café) but not others (e.g. my belief that I'm not a handless brain in a vat). The explanation is that this is because we're often wrong about the former sort of belief, and so we need to track their truth, whereas we are rarely wrong about the latter sort, and so don't need to track their truth.

It is unclear how the analogy is supposed to work. If things had been different, Herring gulls would have had the ability to reliably identify their eggs as well as (or instead of) the ability to reliably identify their chicks. In contrast, even if things had been different, it is hard to see how we could have had the ability to track the truth of beliefs such as that my belief that I'm not a handless brain in a vat. What extra cognitive capacity would make me such that, if I were a handless brain in a vat, I wouldn't believe that I wasn't? This suggests that evolution won't help us make sense of default knowledge. This is, of course, not to say that there isn't a better argument available.

Scepticism

Zalabardo is admirably clear about the extent to which his view vindicates the sceptic. While the sceptic is wrong in thinking that our cognitive activities can never be successful, we can rarely, if ever, know that our cognitive activities have been successful. That Zalabardo doesn't conclude

that his account of knowledge should be rejected in order to avoid this sceptical outcome is perhaps surprising. Zalabardo, in common with many working within the analytical project, sees himself as engaged in a sort of reflective equilibrium. The idea is to get our first-order judgements about which cases are cases of knowledge in line with our general principles concerning the conditions under which someone knows something. One would think that among the relevant first-order judgements are judgements to the effect that we can, at least sometimes, know that our cognitive activities have been successful. Why does Zalabardo discard these first-order judgements, and focus on judgements about Gettier cases and the like?

But perhaps this way of setting things up sees the epistemologist as operating within too narrow a perspective. In Chapter 8 Zalabardo examines the metaphysical presuppositions behind the sceptical problem, and tentatively suggests a way of revising our metaphysics on which his epistemological picture doesn't invite scepticism. While the thought that scepticism raises a problem for our metaphysics rather than our epistemology has a long and distinguished history, it hasn't been quite as prominent in contemporary epistemology. If this thought is back on the table, we shouldn't be too quick to dismiss Zalabardo's account of knowledge. Why abandon an otherwise attractive account because we're working with a false view of the metaphysics?

These issues notwithstanding, *Scepticism and Reliable Belief* is required reading for anyone interested in current thinking about reliabilism and scepticism. The book contains a wealth of resources that any reliabilist can utilise in dealing with standard problems (I found the criticism of Bonjour's 'clairvoyance objection' in Chapter 2 especially convincing). While there may be issues with the details of Zalabardo's particular version of reliabilism, it is a model for how to carry out the analytical project. When reading the book one often gets immersed in the particular puzzles and problems Zalabardo discusses, but he never takes his eye off the bigger project. *Scepticism and Reliable Belief* may not be the end of inquiry into the best form of reliabilism or the prospects for avoiding scepticism, but it makes impressive progress.

ROBIN MCKENNA

Institute for Philosophy
University of Vienna
Universitätsstrasse 7, 1010 Vienna
Austria