

Introduction

It is quite rare in philosophy that one finds consensus, and when one does this tends to provoke a sense of unease. After all, one of the morals that can be gleaned from the study of the history of philosophy is that the emergence of a consensus is often the result of a shared mistake rather than a common understanding. The philosophical consensus that is the focus of this book, and which I argue is most definitely a cause for disquiet, concerns the seemingly universal intuition that *knowledge excludes luck*, or, to put it another way, that the *epistemic luck* that sometimes enables one to have true beliefs (and sometimes even fully justified true beliefs) is incompatible with knowledge. One can see the attraction of the intuition. Knowledge does appear to be a cognitive achievement of some sort, and luck seems to militate against genuine achievements. Accordingly, to say that an agent has knowledge is to imply that she didn't gain that knowledge *simply* via good fortune. Henceforth, we will call this claim that knowledge excludes luck the '*epistemic luck platitude*'.

The problem, however, is that it is difficult to take this thesis entirely at face-value, despite its initial plausibility. After all, there certainly are instances where luck plays a substantial role in the acquisition of apparently bona fide knowledge. Think, for example, of the history of scientific endeavour and those key discoveries that were gained via luck (such as the well-known, though possibly apocryphal, story of Alexander Fleming's supposedly accidental discovery of penicillin), or of familiar tales of plucky detectives who just happen on the vital piece of incriminating evidence at the eleventh hour. Clearly we have here cases—possibly even paradigm cases—in which the agent's true belief constitutes a cognitive achievement even despite the possibly malign influence of luck. Are we to say that the epistemic luck platitude prevents one from ascribing knowledge in these instances?

Nevertheless, whilst there is this tension over just how we are to understand the intuition that knowledge excludes luck, it is clear that there is *something* right about it. Setting aside the issue of whether a guess could ever be an instance

Introduction

of belief for a moment, it is certainly true that your lucky guess that the horse 'Looby-Loo' would win the 4.20 p.m. meeting at Kempton racecourse is hardly knowledge, and at least part of the explanation of why this is so is that it is just plain luck that you were right. Indeed, lucky guesses are an extremal case in this regard, since the point is perfectly general and applies to examples even where no guesswork is involved. The beliefs of the gullible or the epistemically unconscientious, even when true, are often not knowledge, and a good explanation of why this is the case is that the beliefs formed by these agents tend to be only luckily true, and where they are true in such cases it does not represent a cognitive achievement on the part of the agent. We thus need to be careful that in granting knowledge to the Alexander Flemings of this world we do not thereby become so permissive in our knowledge ascriptions as to likewise grant knowledge to those agents whose gullibility or lack of epistemic conscientiousness did not, as it happens, impede them in gaining true beliefs. (And this is not to mention the folly of ascribing knowledge to those epistemically reckless agents who form their true beliefs via lucky guesses.)

Here, then, in essence, is the *prima facie* worry about the claim that knowledge excludes luck. On the one hand, if we construe it too inclusively then we may end up rejecting bona fide cases of knowledge. On the other hand, if we read it too restrictively then we may end up being unable to capture the very sense in which knowledge is a cognitive achievement at all. The issue of how we are to understand the relationship between knowledge and luck is thus central to our understanding of what knowledge is and therefore goes right to the heart of the epistemological project.

The parallel to the more widely discussed debate regarding moral luck is instructive in this regard. Very roughly, the analogy between the two cases goes something like this: just as one cannot take epistemic credit for forming a true belief in a lucky fashion, so one cannot take moral credit for a good consequence of one's action that arose out of luck rather than design (or, for that matter, be morally *blamed* for a bad consequence of one's action that arose out of luck). In short, we are, seemingly, only morally responsible for the consequences of our actions that are in some sense in our control.

Suppose, for example, that by flipping the light switch in my hotel room one evening I inadvertently electrocute the guest in the room next door, with fatal consequences—the result, let us say, of faulty wiring in the hotel that is nothing to do with me, and of which I am (blamelessly) unaware. Clearly, it is just bad luck, in one sense at least, that this happened, and certainly no responsibility of mine. Accordingly, so the intuition runs, I cannot be held morally responsible for this consequence of my actions and thus I am in this respect immune to

Introduction

moral censure. Similarly, the same goes for inadvertent and unforeseeable *good* consequences of my actions, such as if the flipping of this light switch had electrocuted my next-door neighbour *thereby saving his life* (he had, let us suppose, just collapsed from a heart attack, and the burst of current prompted his cold heart back into action). This may be a good consequence that we might loosely describe as being the result of an action of mine, but the luck involved in this good consequence coming about prevents me from taking any moral credit for it.

As with the case of epistemic luck, however, the fact that luck pervades our lives means that there are inevitably going to be cases where understanding moral responsibility in such a way that it *completely* excludes the influence of luck will jar with intuition. For example, consider the case of the otherwise conscientious motorist who just happens to hit a pedestrian at the one moment when his attention is not on the road. It is, we might say in a sympathetic tone, merely a matter of bad luck that this happened and, as in the examples just described, there is certainly a strong case that could be made for thinking that this consequence is both inadvertent and, we can stipulate, unforeseeable (there are usually no pedestrians around this area and, furthermore, the driver had no good reason for thinking that a pedestrian would be around this time). Does this mean, however, that we would be willing to say that this agent is *not* morally responsible for this consequence of his actions (and if we do, would we expect the courts to think likewise)? Clearly our intuitions here are not so straightforward. Thus, as with the phenomenon of epistemic luck, the concern is that we need some way of understanding the claim that moral responsibility excludes luck in such a way that it is neither unduly restrictive nor unduly permissive. And this behoves us to examine in more detail the underlying motivations for this thesis, and the various nuanced ways in which it can be understood.

Although I will be saying something about the problem posed by moral luck, my primary goal in this book is to deal with the far less explored difficulty presented by epistemic luck. Indeed, given the importance of epistemic luck to our understanding of knowledge, it is surprising that there has not yet been a sustained critical discussion of the specific problem posed by epistemic luck. As we will see, a closer examination of this problem highlights some important epistemological truths that have far-reaching ramifications for contemporary epistemology. Furthermore, as I will argue at the end of the book, this discussion of epistemic luck has important implications for the debate about moral luck as well.

In noting that hardly anyone has seriously examined the problem of epistemic luck, I am not claiming that the core thesis in this regard—that

Introduction

knowledge excludes luck—is never *mentioned* (explicitly or otherwise) in the recent epistemological literature, only that it is never really *discussed*. Indeed, as befitting its status as a *universal* intuition—what these days we philosophers tendentiously call a ‘platitide’—one finds this thesis both everywhere and nowhere at the same time. That is, whilst this line of thinking is clearly being presupposed in much of contemporary epistemological thought, the thesis itself is rarely drawn up to the surface of discussion, and even then it is left to stand as it is: a pure platitudinous intuition that is in need of no further explication. For example, Matthias Steup (2001: 2), in a recent survey article on the analysis of knowledge, writes that ‘it is uncontroversial that knowledge is incompatible with epistemic luck’, and simply leaves the matter at that.

My contention here is not that such a claim *is* controversial as it stands, much less that it is straightforwardly false, but rather that we have yet to be given any clear sense of what this claim involves. Given that one core way of understanding the epistemological project is such that it is concerned, in some substantive respect, with providing an account of our key epistemic concepts which explains their incompatibility with epistemic luck, then it is essential that we interpret this supposed platitide correctly. Indeed, it is central to the thesis that I defend in this book that there are various ways of understanding the claim that knowledge excludes luck, not all of them plausible, and that they each pose a different challenge to epistemology.

Consider, for example, some of the areas of epistemological discussion in which the issue of epistemic luck is implicated. One such topic that springs immediately to mind in this regard is the debate concerning the counterexamples to the classical tripartite account of knowledge that were famously presented by Edmund Gettier. These examples work by highlighting the possibility that one might have a true belief that is gained on the basis of excellent grounds (and which therefore meets the rubric set down by the classical tripartite account of knowledge), but which is, nonetheless, and through some strange twist of fate, not knowledge. In effect, what Gettier presented were scenarios in which agents met the conditions specified by the classical tripartite account, but did so in a way that was infected with serendipity.

That just about everyone has been quick to conclude that such examples are successful in their stated aim—in that it is universally acknowledged that the cases described are clearly not instances of knowledge—speaks directly in favour of the claim that knowledge excludes luck. And, indeed, that these examples work because of the incompatibility of luck and knowledge is often explicitly noted. In a relatively recent textbook on contemporary epistemology,

Introduction

for example, we find Jonathan Dancy making the following remark:

justification and knowledge must somehow not depend on coincidence or luck. *This was just the point* of the Gettier counter-examples; nothing in the tripartite definition excluded knowledge by luck. (Dancy 1985: 134, *my italics*)

This is a fairly standard way of understanding the challenge posed by the Gettier counter-examples. Knowledge excludes luck, and thus examples which indicate that one has met the conditions set down by a theory of knowledge, but in such a way that is consistent with a substantive degree of luck being involved, are thereby conclusive counter-examples to that theory of knowledge. Accordingly, it should come as no surprise to learn that it is hard to find anyone these days who would seriously defend the classical tripartite theory of knowledge.

This neat picture of the relationship between knowledge and luck is complicated, however, once one considers how the issue of epistemic luck figures in other epistemological debates. A good contrast in this regard is the debate regarding radical scepticism. Again, we often get the intuition expressed that part of the sceptical challenge is to offer an analysis of knowledge that does not allow that our putative 'knowledge' is susceptible, in a wholesale manner, to luck. Just such an intuition seems to be behind, for example, infallibility-based approaches to scepticism where the leading idea seems to be that our knowledge must be infallible because otherwise it would be subject to the vagaries of luck.

Nevertheless, it is also interesting to note that in the debate regarding scepticism a certain degree of *acceptance* of some form of epistemic luck is often countenanced as well. One finds, for example, Michael Williams arguing as follows:

Knowledge and justification always involve an element of *epistemic luck*. A belief whose truth is *wholly* accidental cannot count as knowledge. But getting things right is never wholly nonaccidental either. (Williams 1999: 59, *italics in original*)

Similarly, Linda Zagzebski (1999: 109) has recently written that 'epistemic luck permeates the human condition, whether for good or ill'. Without any specification of what is involved in epistemic luck, however, it is far from clear what these concessions, if that's what they are, amount to. Moreover, neither is it clear just how allowing some degree of epistemic luck in this case is meant to be compatible with apparently not allowing *any* degree of epistemic luck in the Gettier cases.

The principal goal of this book is to capture what is involved in epistemic luck, and thus to identify the sense in which luck is incompatible with

Introduction

knowledge and the sense (if any) in which it isn't. Since one finds the most explicit discussions of epistemic luck in the debate regarding radical scepticism, the first part of the book focuses on this debate. (Indeed, as we will see in Part II, the fundamental problem posed by epistemic luck is a sceptical problem.)

In Chapter 1 I examine the way in which scepticism is often motivated via an appeal to some kind of infallibility thesis—as expressed, for example, in the work of Peter Unger—and highlight how this construal of the sceptical problem presupposes a certain view about epistemic luck. Furthermore, I explore the manner in which an infallibilism-based scepticism gives rise to the kind of template sceptical argument that is the focus for much of the contemporary discussion of scepticism and which turns on the principle that knowledge is 'closed' under known entailments (the so-called 'closure' principle).

In Chapter 2 I critically examine two of the main responses to the closure-based sceptical argument. In particular, I look at the case that has been made for the rejection of closure by, for example, Fred Dretske and Robert Nozick, and examine the so-called 'sensitivity' condition on knowledge that they advocate in this respect. I then evaluate this approach in the light of the attributer contextualist response to scepticism, as developed by such figures as Stewart Cohen, Keith DeRose, and David Lewis. In both cases I argue that these antisceptical approaches suffer by being motivated by diverse and incompatible intuitions, principally as regards their ambivalent endorsement of epistemological externalism. An exploration of the relationship between the sceptical problem and the epistemic externalism–internalism distinction is thus presented.

In Chapter 3 I contend that in so far as the contemporary debate about scepticism has formulated the sceptical problem correctly, and in so far as one is entitled to adopt epistemological externalism as part of one's antisceptical strategy, then one should reject both the arguments for non-closure and for attributer contextualism, and advance instead the kind of neo-Moorean position—as defended in the recent literature by, amongst others, Ernest Sosa and myself—which turns on the so-called 'safety' condition on knowledge. This chapter is thus devoted to exploring the implications of this approach for the sceptical problem, including the manner in which it needs to be distinguished from the more familiar 'Moorean' approach to scepticism that is associated with G. E. Moore's famous proof of an external world. A key component of this chapter in this respect is an account of the conversational rules that govern our ascriptions of knowledge (including self-ascriptions) which can do justice to the intuitions that motivate attributer contextualism whilst also explaining

Introduction

why the kind of antisceptical assertions made by Moore are so problematic. As regards the latter, I enlist elements of Wittgenstein's critique of Moore's response to the sceptic from *On Certainty*. I conclude that what these considerations highlight is a deep and important (and also surprising) truth about the structure of reasons.

Crucially, however, I do not pretend (as others do) that the neo-Moorean approach is able, by itself at any rate, to resolve the sceptical problem. As I explain in Chapter 4, the reason for this is that the contemporary discussion has misunderstood what the focus of scepticism is, and thus even the most plausible of the three main antisceptical theories in the literature misses its intended target. In order to illustrate this, I contrast the closure-based template sceptical argument with its analogue argument expressed in terms of the 'underdetermination' principle that can be found in ancient Pyrrhonian sceptical writings and which merely demands that one's evidence should favour one's beliefs over incompatible alternatives. I show how the latter argument brings to the fore just what is wrong with the contemporary debate regarding scepticism, which is its failure to respond to the specifically internalist and evidentialist character of the sceptical doubt.

At this point the reader might begin to wonder—entirely legitimately I think—just what all this has to do with epistemic luck. For whilst, as noted above, there is a clear sense in which the sceptical problem seems to trade on intuitions we have about the relationship between luck and knowledge, it is not immediately obvious why a deeper analysis of epistemic luck should cast any light on the sceptical debate. As I show in Part II, however, the confused understanding of the sceptical problem in the contemporary epistemological debate is indicative of a deeper confusion regarding the status of the claim that knowledge excludes luck. The reader is thus asked for patience while I set out the main responses to scepticism in the literature, since the extensive groundwork of Part I is necessary if the detailed discussions of epistemic luck in Part II are to be possible.

Indeed, there is a second sense in which the groundwork on scepticism undertaken in Part I is necessary for the specific discussion on epistemic luck that takes place in Part II, and this concerns how it is in response to the problem of scepticism that most of the main currents of contemporary epistemology have been motivated. In particular, sensitivity-based, safety-based, and attributer contextualist theories of knowledge all started out as antisceptical theories but are now theories of knowledge motivated on grounds that are independent of the problem of radical scepticism. Identifying a problem regarding epistemic luck that lies at the heart of the debate about radical scepticism

Introduction

is thus tantamount to discovering that contemporary epistemology as a whole is in trouble. The extended critique of Part I thus introduces all the main positions and thereby sets the scene for the second part of the book where critique starts to give way in favour of reconstruction.

I begin Part II with an overview of the various claims that have been made concerning epistemic luck and present an elucidation of the concept of luck in the light of this discussion that, I argue, captures the essence of the notion. With this elucidation in mind, I then identify—with the help of Unger—several varieties of luck that might be referred to as epistemic and which are entirely compatible with knowledge possession.

With these benign types of epistemic luck set to one side, I then turn, in Chapter 6, to the two species that remain and explain how they feature in the main epistemological debates. The first species of epistemic luck—what I call ‘veritic’ luck—can be handled with a modest ‘relevant alternatives’ account of knowledge that is specifically defined so that it counters this type of epistemic luck. As I explain, such a theory is essentially a version of the safety-based neo-Moorean thesis that we looked at in Chapter 3. I develop this view by considering some of the main examples that are discussed in epistemology—such as Gettier counter-examples, the ‘lottery’ puzzle, and ‘barn façade’ examples—and show how the neo-Moorean account can in each case generate the right result. Along the way the formulation of this anti-veritic-luck thesis is refined in response to potential challenges that might be presented to the view.

Significantly, however, the chapter closes by arguing that there is a second species of epistemic luck—what I call ‘reflective luck’—that is epistemically problematic and which is not dealt with by the neo-Moorean account. The bulk of the rest of the book is devoted to getting to grips with this type of epistemic luck. In particular, in Chapter 7, I argue that such a notion raises important questions about the centrality to epistemology of a conception of justification which demands that one is able to take cognitive responsibility for one’s beliefs. I also look (critically, as it turns out) at some of the recent ‘virtue epistemologies’ that have been put forward in the recent literature which (in essence) define knowledge in terms of the epistemic virtues (and also, in some cases, the cognitive faculties).

More specifically, I contrast broadly externalist construals of the virtue-theoretic thesis—such as the ‘agent reliabilism’ advanced, for example, by Sosa and John Greco—with virtue epistemologies that roughly fall into the internalist camp—in particular, ‘neo-Aristotelian’ theories, such as the view defended by Zagzebski. I claim that in each case a key part of the motivation for offering such a view lies in a desire to eliminate a species of epistemic luck that

Introduction

is thought to be left uneliminated by rival theories of knowledge. In the case of agent reliabilism, the rival view is the similarly externalist ‘process reliabilist’ theory of knowledge advanced by, for example, Alvin Goldman. For the neo-Aristotelians, in contrast, the rival views are externalist theories of knowledge more generally, including agent reliabilism itself. Crucially, however, I maintain that whilst the kind of epistemic luck that Greco and company are focusing upon is specifically veritic, those offering a Zagzebski-style virtue theory are, it turns out, concerned also with the exclusion of reflective luck. In general, I claim that externalists and internalists in epistemology often speak past one another precisely because they are concerned with these different types of epistemic luck, and that disambiguating the species of epistemic luck at issue in this debate can help us to find a way of reconciling both internalist and externalist intuitions about knowledge.

More dramatically still, I conclude this chapter by arguing that once the ambiguous role of epistemic luck in the development of these virtue epistemological theories is made explicit, then the motivation to offer a specifically *virtue*-theoretic theory of knowledge subsides. In particular, I show how, pending further detail about the virtue-theoretic account at least, one can achieve everything that virtue epistemologists claim to achieve with their theories without adopting the core virtue-theoretic claim that knowledge must be defined in terms of the epistemic virtues (and cognitive faculties). Accordingly, there is no need to endorse the more radical theory of knowledge proposed by virtue theorists. A close analysis of epistemic luck thus not only helps us to see a way through the apparently intractable debate between epistemic externalists and internalists, but also appears to undercut the motivation for one of the most influential epistemological proposals of recent years.

In Chapter 8 I return to the sceptical challenge in the light of this distinction between veritic and reflective epistemic luck and argue that the inadequacy of the antisceptical proposals considered in Part I is a result of how they only (at best) eliminate veritic luck, and thus do not engage with the problem of reflective luck at all. Crucially, however, I claim that it is the specific challenge posed by reflective luck that is central to the sceptical problem, and yet that there is a fundamental sense in which this type of epistemic luck is ineliminable. In order to add some content to this claim, I argue that it is this sceptical problem that informs the Pyrrhonian sceptical challenge of antiquity. Moreover, I further maintain that the so-called ‘metaepistemological’ sceptical challenge that features prominently in contemporary epistemological debate—as advanced, for example, by Barry Stroud and Richard Fumerton—is best understood in terms of the specific sceptical problem regarding the ineliminability of reflective luck.

Introduction

Chapter 9 presents a more in-depth discussion of the difficulty posed by the ineliminability of reflective luck and offers one way of responding to the difficulty. I begin by reconsidering Wittgenstein's remarks on knowledge from *On Certainty* that we looked at in Chapter 3, and consider in particular the specific claims he makes about the so-called 'hinge' propositions that contextually determine the nature of epistemic evaluation. I claim that what underlies Wittgenstein's remarks in this respect is a certain view about the ultimately groundless nature of our reason-giving practices. This way of thinking about the structure of reasons is, I argue, correlative to the thesis described here regarding the ineliminability of reflective luck. Moreover, I contend that it is our implicit philosophical recognition of this problem that gives rise to a certain kind of general anxiety about our epistemic position that I refer to as epistemic *angst*. This contention is further illustrated by critically evaluating John McDowell's content externalist response to the sceptic which incorporates a strongly antisceptical account of reasons in the light of the opposing Wittgensteinian scepticism-friendly account.

This final chapter of the main body of the book is not completely pessimistic, however, in that it closes by arguing that whilst there is no epistemic response available to the problem posed by reflective luck, there is a plausible *pragmatic* resolution to this difficulty. Indeed, I claim that an antisceptical thesis of this sort can be discerned from Wittgenstein's own fragmentary remarks on scepticism, and I develop such a view in the light of Hans Reichenbach's pragmatic response to the problem of inductive scepticism, which I maintain is cast along similar lines.

The book closes with a postscript which outlines how this analysis of epistemic luck has ramifications for the supposedly parallel debate regarding moral luck. As noted above, both discussions seem to rest upon the assumption that the moral and epistemic states in question are (in some sense) achievements and so cannot be subject to luck. Nevertheless, and as we also noted above, whilst tacit acceptance of this brute intuition is common in contemporary epistemology, contemporary moral theory has been conspicuous in its sustained scrutiny of this claim. Focusing on the two classic papers in this area by Thomas Nagel and Bernard Williams that sparked the modern debate on this subject, I examine the arguments and examples put forward in this regard and highlight how the primarily epistemological analysis offered here can be put into service to cast light on the corresponding moral debate. In particular, I show that these authors fail to distinguish a specifically moral problem, and that what difficulties they do highlight are parasitic on the problem of reflective epistemic luck that is identified here. I therefore conclude that there is a strong

Introduction

prima facie case for thinking that there is no such thing as moral luck, and that the only reason why many think that there is is because they are confusing the non-existent problem of moral luck with the entirely genuine problem of reflective epistemic luck. Our discussion of epistemic luck thus has repercussions which extend beyond the purely epistemological realm.

