

# The Standard Analytic Conception of Knowledge

In spite of the multitude of energetic epistemological debates presently occurring, epistemology's core maintains a deep contentment with several theses about the nature of knowledge. Individually, these are rarely questioned. Collectively, they constitute a partial conception of knowledge. It is a partial conception that is reflective of much about what is often called analytic epistemology, insofar as such epistemology talks about knowledge at all.<sup>1</sup> But it also deserves not to be so readily accepted by so many philosophers. This chapter will outline that partial conception in a generic way, before indicating in some equally generic ways why it deserves to be questioned, even modified, possibly modified significantly.

## 1.1 'Knowing is a Belief State (or Something Similar)'

To know is to be in a particular state; or so it is routinely assumed by epistemologists. And that assumption has the following implications. Knowing that  $p$  and knowing that  $q$  are different particular states. A typical knower is thus in many particular states at once, as she knows that  $p$ , knows that  $q$ , knows that  $r$ , and so on. And her state of being a knower in general is some function of her being in those many particular states of knowing.

In understanding what it is to be a knower at all, then, we may focus on some arbitrary one of those alleged states — that of knowing that  $p$  (for an

<sup>1</sup> As Riggs (2008) reminds us, there can be much more to epistemology than a conceptual focus upon knowledge and its details. My epistemological concern in this book *will* be just with knowledge and its details, though.

## 2 *The Standard Analytic Conception of Knowledge*

arbitrary ‘ $p$ ’). What kind of state is it? Most epistemologists have long favoured this answer to that question:

Knowledge is a kind of belief. To know that  $p$  is, at the very least, to be in a state of believing that  $p$ .

In one forum after another, epistemologists assure us that knowledge is — indeed, that it must be — a suitably enhanced or impressive belief.<sup>2</sup> The various forms of enhancement or impressiveness that are claimed by different epistemologists to be required will be gestured at in Section 1.2, but the immediately pressing point is the near-unanimity among epistemologists as to which aspect of a person it is that needs to be impressive if knowledge is to be present. Beliefs do not just report our knowledge; they *are* our knowledge. Admittedly, only some of them are; no non-beliefs are, though.

Sometimes, variants of that idea are proposed. For example, Keith Lehrer (1990: 10–11) argues that knowledge is always a kind of acceptance, while allowing that this can be said to be a special kind of belief.<sup>3</sup> Laurence Bonjour (2002: 30) allows knowledge to be belief or acceptance; as — more complicat- edly — does Jonathan Cohen (1992: ch. IV). And Ernest Sosa (1980: 3) embraces what is potentially an even wider categorisation of the state of knowing that  $p$ . He says that ‘nothing can be known without being at least believed (or accepted, or presumed, taken for granted, or the like) in some broad sense.’

In one way or another, then, most epistemologists accept either the *knowledge-as-belief* thesis or the *knowledge-as-belief-or-acceptance-or-something-similar* thesis. Each thesis is regarded as a special case of this *knowledge-as-state* thesis:

Knowledge is a kind of state. To know that  $p$  is, at the very least, to be in some kind of  $p$ -directed state.

That is a highly schematic thesis. How do we render it less so? Acceptance of the knowledge-as-belief thesis has been the paradigmatic means by which epistemologists claim to understand the knowledge-as-state thesis.

That said, though, the knowledge-as-state thesis has remained at the centre of epistemological thinking even when, as occasionally occurs, a philosopher

<sup>2</sup> For just a few (chosen almost at random) from the vast number of examples of this classification, see Ackerman (1972: 71–3), Chisholm (1989), Moser (1989: 13–23), Zagzebski (1999: 92–3), Audi (2003), and Feldman (2003: 13–14). Strictly, epistemologists often say just that knowing *entails* believing. But (it is assumed) something needs to *be* the knowledge, even if only to be what has the further entailed properties. And the belief that  $p$  that is taken to be entailed by the knowledge that  $p$  is standardly presumed to play this role.

<sup>3</sup> Acceptance is deemed by Lehrer to be ‘aimed at truth’, whereas not all beliefs are like that (1990: 11).

seeks to *avoid* analysing knowledge as a belief (or acceptance, or anything similar). The most notable instance of this avoidance was Plato's, in the *Republic* (476d–480a). He distinguished between the state of knowing and the state of believing, without proceeding to reduce the former to the latter, or indeed to anything comparable. These states were to be distinguished, most obviously, in terms of their objects. Knowledge (said Plato-of-the-*Republic*) is a state in which one is related to what *definitely* is, what *cannot* be other than what is — in short, a necessary truth. In contrast, belief is a state of being related to what-*is-yet-need-not-be* — in short, a contingent truth. The distinction also points to knowledge being a state that arises from the exercise of reason, while belief is brought about by perception.

So said Plato; and some recent thinking has overlapped here with Plato's. Zeno Vendler (1972: ch. 5), Kenneth Sayre (1997: ch. 5), and Timothy Williamson (2000) regard knowing as a state that is not explicable as a state of belief. Indeed, for Sayre and for Williamson, knowing is a *primitive* state. They deny that it is reducible to, or analysable as, *any* further sort of state. Although (says Williamson) knowing 'is a state of mind' (2000: 21), it 'does not factorise as standard analyses require' (2000: 33). Rather, 'knowing is the most general factive stative attitude, that which one has to a proposition if one has any factive stative attitude to it at all' (2000: 34). And Sayre (1997: 139 n. 9) contends

that knowledge is a cognitive state that cannot be analyzed into more basic cognitive components (such as evidence and belief). The present approach [by Sayre] agrees with Plato's treating knowledge as cognitively basic.

Sayre and Williamson are thus heirs to an ancient urge. Even while rejecting the dominant contemporary view of knowledge as a kind of belief, they retain the more general, but epistemologically no less standard, commitment that usually underlies that contemporary view. That is, they accept the thesis that knowledge — factual or propositional knowledge, knowledge that  $p$ <sup>4</sup> — is some kind of state of the knower: An individual's knowing that  $p$  is her being in an appropriate state. This, it seems to epistemologists as a whole, is an unquestionably true thesis about knowledge.

And perhaps it *is* true, when formulated so generically. But epistemologists do not accept only so generic a thesis. In various ways, they accept that thesis *by* accepting instances of this comparatively generic thesis:

Knowledge is a state — either of belief or something similar, or primitively or unanalysably so.

<sup>4</sup> I use both of the terms 'factual' and 'propositional' here, so as not to beg the metaphysical question of whether knowledge that  $p$  is knowledge of a fact or instead knowledge of a true proposition. For a semantic treatment of this metaphysical issue, see Moffett (2003).

#### 4 *The Standard Analytic Conception of Knowledge*

That is, an accurate conceptual analysis of knowledge — if even possible — would deem it to be a belief or something similar. (So, I call this the *knowledge-as-either-a-belief-or-an-unanalysable-state* thesis.)

### 1.2 ‘Knowledge is Well Supported’

Epistemologists are no less committed to the thesis that, whatever else knowledge is, it is something that is somehow suitably enhanced or impressive. Standardly (we are told), there are two aspects to such enhancement or impressiveness:

- (1) Knowledge involves factivity. Insofar as knowledge is a state, the state is factive. Insofar as knowledge is a belief, the belief is true.
- (2) Nothing — no state, such as a belief — is knowledge if it is not somehow well supported.

I will not be questioning (1);<sup>5</sup> we need to be clear on what (2) means, though. I have used the generic term ‘well supported’ with the intention of encompassing, as neutrally as possible, the multitude of more-or-less-specific suggestions that have been made on this issue. ‘Justified’ is the term most commonly used in this connection.<sup>6</sup> But some philosophers (such as Lewis 1996: 551) have used the term ‘justified’ more narrowly, affixing it only to a person’s evidence or reasons; while others have reserved a term such as ‘epistemized’ (Alston 1989: 58) or ‘warranted’ (Plantinga 1993a: 3) for whatever enhancement most clearly distinguishes knowledge from ‘mere’ true belief. No matter: the traditional *knowledge-as-well-supported* thesis is broad enough to absorb these various approaches. It says only that knowledge requires the presence of some feature — something suitably impressive — beyond belief (or beyond whatever else plays belief’s role within the knowledge) and truth.

What is that ‘something’? Maybe the presence of an internally coherent body of evidence is the key (BonJour 1985). Maybe what is needed is evidence that does not overlook too many crucial aspects of the believer’s neighbourhood (Lehrer and Paxson 1969). Maybe it is enough if the pertinent belief’s genesis is sufficiently reliable (Goldman 1979); or maybe the belief has to be functioning aptly in its environment (Plantinga 1993b). And so forth. There have been myriad such suggestions (and detailed discussion

<sup>5</sup> But Hazlett (2010) does so.

<sup>6</sup> For a few of the unmanageably large number of epistemological endorsements of this generic idea, see Moser (1989: 35–7), Lehrer (1990: 12–3), BonJour (2002: 38–43), Audi (2003: 4), and Feldman (2003: 15–6).

of them could reasonably occupy a philosophical career). What unites these otherwise disparate epistemological efforts is a commitment to this knowledge-as-well-supported thesis:

Nothing is knowledge if it is not well supported. For example, no belief, even if true, is knowledge unless it is well supported.

That thesis is usually stated in these simpler and more specific terms:

No merely true belief is knowledge; some suitable enhancement is also needed.

For simplicity, I will often focus upon that more specific thesis.

Note that it is a thesis, as is the more generic one, that is standardly advanced by epistemologists as being something more than a merely contingent and empirically supported truth. Epistemologists do not mean to say that, although all of this world's instances of knowledge happen to be well supported, things could have been otherwise in this regard. On the contrary: epistemologists embrace more or less generic versions of the knowledge-as-well-supported thesis with remarkable confidence, a degree of confidence more congruent with regarding the thesis as a necessary truth than as a contingent one. In either the same or a related spirit, what is often said by epistemologists is that the thesis is a conceptual truth. They make claims to this effect: 'It is part of the very concept of knowledge that no merely true belief, unaccompanied by good support, is knowledge.'

And, if the question of that putative truth's epistemic standing were to arise (as may well occur, especially in philosophical contexts), undoubtedly epistemologists would claim to *know* that knowledge is as the knowledge-as-well-supported thesis describes it as being. Moreover, they would not do so only by adverting to some specific and technical theory of knowledge, with this being what legitimates their claim to know the thesis to be true. Rather, they would be more likely to describe that thesis — that putatively conceptual truth — as being known via only a little reflection, without calling upon some detailed theory of knowledge in support of this description. The thesis would be said to be manifestly true or intuitively true, for example.

Nor is such confidence in the truth of that thesis newly arrived within philosophy. Even Socrates, the master disavower of knowledge, laid claim to this particular piece of knowledge. In the *Meno* (98b), he announced that he knew this principle — what I am calling the knowledge-as-well-supported thesis — to be true, even while he continued to disclaim almost all other knowledge.<sup>7</sup> He averred, with unusual

<sup>7</sup> I follow Benson (2000: 8–10), who was following Vlastos, in accepting that the early Platonic dialogues represent the views of Socrates and of the 'early' Plato, whereas the middle

confidence, that knowledge is more than a true belief: even a true belief is knowledge only when suitably enhanced.

We are in the presence, therefore, of an exceptionally long-lived and central commitment within epistemology. Epistemologists will generally insist that, if they know anything at all about knowledge that *p*, they know that it involves — indeed, that it has to involve — some form of good support, such as would be provided by good evidence for the truth of *p*.<sup>8</sup>

### 1.3 ‘Knowledge is Absolute’

The epistemological commitments outlined in Section 1.1 and Section 1.2 are almost always explicit within any given epistemologist’s writing. But not all epistemological commitments are present so overtly; some only silently impart structure to epistemological thinking. One commitment whose presence within most epistemological thinking is implicit, rather than explicit, is a thesis of *knowledge-absolutism*. This is the thesis that knowledge — specifically, knowledge that *p* — is absolute. According to this thesis, no knowledge of a particular truth ever admits of varying *grades* (either within a particular context or across different contexts).<sup>9</sup>

Knowledge-absolutism thus implies that there cannot be two instances of knowledge that *p*, one of which is somehow a better or higher grade of knowledge that *p* than is the other. So, in particular, no instance of knowledge that *p* is ever improvable purely as knowledge that *p*.<sup>10</sup> Once a belief is

---

and late dialogues reflect the more distinctively Platonic theorising. The *Meno* instantiates the former category, whereas the *Republic* is an instance of the latter. Still, Benson does acknowledge (2000: 94) that the *Meno*’s distinguishing of knowledge from mere true belief ‘is correctly thought to represent a new emerging Platonic view’.

<sup>8</sup> A few epistemologists do dissent from this consensus, as Chapter 4 will note.

<sup>9</sup> Occasionally, epistemological commitment to this thesis is made explicit. Recent examples include Stanley (2005: ch. 2), Bird (2007: 106–7), Elgin (2007: 36), and Sutton (2007: 153 n. 42, 154 n. 4). Section 5.12 will evaluate Stanley’s arguments. Elsewhere (2001a: 13–8), I have discussed two earlier instances of philosophers — Ryle (1949: 54–5, 59) and Dretske (1981: 363) — accepting knowledge-absolutism. Still, not all epistemologists accept that thesis. In Section 2.7 I discuss several who do not. (And in Section 2.8 I comment on a few, including contextualists, whom we could readily *but mistakenly* believe do not.)

<sup>10</sup> Hence, too, knowledge-absolutism is not the denial of there being better or worse *kinds* of knowledge — such as if knowledge that *p* is claimed to be epistemically superior to knowledge that *q*, because *p* is a truth from a domain that admits of being better known than do truths, including *q*, from some other domain. The history of philosophy includes that idea — for instance, with pure reason being thought to provide a higher quality of knowledge than do the senses. But the possibility I will investigate is of there being different grades of knowledge within a *single* domain, indeed of a *single truth*. See, for example, Jacques Maritain (1959). See, too, Wuellner (1966: 164) on ‘degrees (modes) of knowledge’. I discuss these ideas more fully in Section 2.6, Section 2.7, and Section 2.8.

sufficiently well supported (all else being equal) to be knowledge that  $p$ , it cannot become better purely as knowledge that  $p$ . Not even by becoming better supported could it improve *qua* knowledge that  $p$ . For example, a more evidentially justified instance of knowledge that  $p$  is not better as knowledge that  $p$  than is a less evidentially justified instance of knowledge that  $p$ . The better evidentially justified instance might be preferred for independently good reasons, such as when the extra evidence provides appropriately increased confidence in the truth of the belief that  $p$ . Better evidence can have its own benefits. And there can be more or less of these, as — more generally — there can be better or worse support for the truth of a belief. Even so, the absolutist view of knowledge, which is common among epistemologists, has been that, once enough support is present to make a belief that  $p$  knowledge (other things being equal), knowledge that  $p$  is present — until, for whatever reason, it is no longer present. And that is that. The knowledge cannot fluctuate in quality as knowledge. It can only be — or not be. *Qua* knowledge, it can only be present or absent. It cannot be more or less present or absent — even as, all the while, it is present.

An analogy may clarify this characterisation of knowledge-absolutism. In baseball, home runs can vary in qualitative properties, such as the amount by which they clear the outfield wall or the speed at which they travel. But none of that variability affects the fact of whether or not a home run has been struck. And knowledge — according to knowledge-absolutism — is like that. Even if the ball only just fails to clear the outfield wall, there is no home run. Even if it only barely clears the fence, there is a home run. And once there is one, nothing else about the quality of the home run — such as its power or extra length — constitutes how, *officially*, it is a home run. Although a commentator may remark upon those variations in quality, they contribute nothing to the home run's *officially being* a home run.<sup>11</sup> Analogously, knowledge-absolutism tells us that no instance of knowledge that  $p$  is *better purely as* knowledge that  $p$  than a second instance is, even if the first one is better than the second one in respects that happen to be part of the two being the respective instances they are. Home run 1 is no better than home run 2 purely as an official home run — even if in fact home run 1 has been hit with greater force, thereby clearing the outfield fence more easily, than home run 2 has been. Equally, instance 1 of knowledge that  $p$  is no better purely at being knowledge that  $p$  than is instance 2 of knowledge that  $p$  — even if instance 1 is supported by more good evidence, say, than instance 2 is (given that each is well-enough supported to be knowledge that  $p$ ). Knowledge-absolutism thus denies that there could be a qualitative hierarchy among instances of knowledge of a particular  $p$  — insofar purely as each of these *is* knowledge that  $p$ . This denial is maintained even while

<sup>11</sup> My uses here of '*officially*' correspond to my use of '*qua*' in the previous paragraph.

allowing that the different instances of knowledge that *p* could be arranged hierarchically in other respects, such as the respective strengths of the bodies of evidence contributing to their being the particular instances they are of knowledge that *p*. (Instance 1 of knowledge that *p* could be based upon better evidence than instance 2 is — without thereby being better purely *as* knowledge that *p*.)

Why has this view of knowledge taken hold among present-day philosophers? Here, I note only that it is a picture of knowledge that might strike so-inclined epistemologists — those looking to support knowledge-absolutism — as combining fluently with Section 1.1's knowledge-as-either-a-belief-or-an-unanalysable-state thesis. Knowledge would be thought of as a state that is either present or absent, while being unable to be more or less present or absent. There are different possible models for that sort of state. To take just one: Knowledge that *p* would be akin to a state of being adequately illuminated — a state that is wholly present or wholly absent, in each case at the flick of a switch.

Section 1.2's knowledge-as-well-supported thesis, too, will readily be thought by most epistemologists to contribute to our understanding (as follows) of that kind of picture. At one moment, without sufficient support being present, knowledge that *p* is not present. At the next moment, with sufficient support becoming present (and with all else being equal), knowledge as such comes into existence, with the crucial threshold of support having been reached. In that sense, knowing may continue to be thought of as like a room's being properly illuminated. The relevant switch is flicked; suddenly, the room is properly illuminated — only now, though, not previously. And, we are standardly being told, knowledge is like that. No matter how close a situation has previously been to containing knowledge that *p*, it comes to include that knowledge only once the crucial threshold of support is reached. Moreover, once that threshold *qua* threshold has been reached, it cannot be reached even more (so to speak) *qua* threshold. The room is properly illuminated; or it is not. The aptly installed light bulb is working; or it is not.

Again, it seems that most epistemologists routinely regard knowledge that *p* as being like that. There is sufficient illumination, amounting to knowledge of a particular truth — or there is not. There is knowledge only insofar as there is sufficient illumination; and the latter is all or nothing. Once there is sufficient illumination, there cannot be more-than-sufficient illumination; the room is already sufficiently illuminated.<sup>12</sup> Section 1.2 implied that only once enough support is present is there a state of knowledge that *p*; the

<sup>12</sup> 'But a light may be further controlled, such as by a dimmer switch — one that dims or intensifies the illumination within the room.' Sometimes that does occur. Can knowledge that *p*, analogously, be like that? To think so is to doubt knowledge-absolutism (as indeed Chapter 2 and Chapter 5 will do).



present section adds to that implication the claim that, once such a state is present, there cannot be any more to the presence of that state as such. That state is that state is that state — nothing less, but also nothing more.

## 1.4 ‘Knowing Includes not being Gettiered’

Contemporary epistemologists in general accept the knowledge-as-either-a-belief-or-an-unanalysable-state thesis (Section 1.1), the knowledge-as-well-supported thesis (Section 1.2), and knowledge-absolutism (Section 1.3). That list also includes this *knowledge-as-not-Gettiered* thesis:

Knowledge is present whenever some concatenation of circumstances or features is, only if that concatenation is not *Gettiered*. (That is, the concatenation is not present as part of a situation of the kind that has come to be called a ‘Gettier case’.)

A significant amount of epistemology since 1963 has identified, then tried to solve, what is generally called the Gettier problem.<sup>13</sup> Often, this has involved epistemologists trying valiantly, yet inconclusively, to show why it is true that having knowledge is not so undemanding as ever to be attainable by a person’s having a belief that is only luckily both justified and true. Other characterisations, too, of Gettier cases have been hypothesised and investigated; and Chapter 3 will discuss how to describe such cases. The immediate point, however, is simply that almost no epistemologist believes that something can be Gettiered *and* knowledge.

The Gettier problem could be thought of as a special case of the problem of ascertaining what kind of justification is needed within knowledge. But because epistemologists tend not to treat the Gettier problem in that way, I will retain a distinction between these two issues. Accordingly, they may be thought of as *two* aspects of knowledge’s core. There would be a justificatory-core and a Gettier-core. Epistemologists are routinely adamant that to specify accurately the amount and kind of justification minimally needed within knowledge is to describe what is definitively part of knowing. They are no less adamant about the same being true of any definitive specification of how to evade the Gettier problem. At this stage of the book, we need only accept that, according to epistemological orthodoxy, there *is* such a problem, resolution of which is required if part of knowing’s core is to be espied. *Gettierism* is what I will call this orthodoxy.

<sup>13</sup> It stems from Gettier (1963) and will be Chapter 3’s topic. For more on the putative problem’s history, see Shope (1983) and Hetherington (1996; 2005*b*; 2010*b*).

## 1.5 ‘Knowledge-that is Fundamentally Theoretical, not Knowledge-how’

Another element of knowledge’s putative core is categorial. As a matter of professional history, epistemologists have long sought to understand propositional knowledge in particular, when trying to understand knowledge at all. Indeed, epistemological discussions generally use the word ‘knowledge’ purely to designate propositional knowledge (even if not in principle, at least for convenience). And there is a time-honoured reason for that fixity of professional gaze. Epistemologists have routinely trusted that, if they can understand propositional knowledge’s nature, then (1) they will have uncovered the nature of the specific form of knowledge supposedly most central to human inquiry, especially to scientific inquiry, and (2) they could thereby be well-positioned to analyse, in turn, what may seem to be other kinds of knowledge (such as by conceptually reducing these to some version or function of propositional knowledge, thus understood).<sup>14</sup>

Let us remind ourselves of what these *prima facie* ‘other kinds of knowledge’ are. We talk of knowing a place (‘I have known this town for years, worse luck!’); we claim to know how to perform a task (‘Fortunately, I do know how to ride a bicycle; I can therefore save the world’); we believe that we know other people (‘Yes, yes, I admit that I know him’); and so on.

One of these in particular — knowledge-how, knowledge of how to do something, practical knowledge — has often been thought to be notably different in kind, categorially distinct, from propositional knowledge.

Indeed, when initially laying out their subject’s *explananda* (those phenomena requiring explanation), epistemologists standardly assume from the outset that knowledge-how is to be distinguished fundamentally from knowledge-that. Occasionally, an attempt is made to question this, by showing that, and how, knowledge-how is really a kind of knowledge-that. Even then, though, knowledge-that remains on its conceptual throne. It remains a theoretical kind of knowledge.

Chapter 2 will discuss all of this — the claim of categorial distinctness, and the idea of knowledge-how being a kind of knowledge-that — in more detail. The present section’s point is merely that the following *knowledge-that-as-fundamentally-theoretical-knowledge* thesis — which

<sup>14</sup> Of course, sometimes epistemologists discuss only propositional knowledge, ignoring non-propositional knowledge, simply because, as Zagzebski (2009: 5) puts it,

(1) It is very difficult to analyze it [i.e., non-propositional knowledge] and it is hard to say anything about it that adds to our understanding of it, and (2) It is so different from propositional knowledge that it needs a separate treatment.

we may call a *theoreticalism* about knowledge — has a secure place within epistemological orthodoxy:

Knowledge-that is fundamentally theoretical knowledge — in the sense that it is not knowledge-how (practical knowledge). Either knowledge-how is a categorially distinct kind of knowledge from knowledge-that, or it is best understood in terms of knowledge-that (which is not itself to be understood in terms of knowledge-how).

## 1.6 The Standard Analytic Conception of Knowledge

Now let us combine Section 1.1, Section 1.2, Section 1.3, Section 1.4, and Section 1.5. The result should be readily recognisable. So much so that I will refer to it as ‘the’ standard analytic conception of knowledge.

In doing so, I do not mean to insist that every epistemologist accepts each of the conditions identified in those sections. Still, almost all epistemologists accept either all or almost all of those conditions. Nor do I mean to deny that many individual epistemologists will also offer favoured further conditions (sometimes, these being different conditions for different individual epistemologists). No matter; epistemologists’ doing this is consistent with my claim that the earlier sections have jointly given us the core of a conception that guides much epistemological research within what is usually called ‘analytic philosophy’. This is the conception I have in mind:

Knowledge-that is a state (either unanalysably so or, for instance, a belief). At base, it is theoretical knowledge; it is not a kind of knowledge-how. It is well supported (thus, not merely a true belief). It is absolute, unable to admit of differing grades. And it is not Gettiered (whatever, precisely, this turns out to be).

Then we may summarise that description:

Knowledge-that is impressively and absolutely theoretical knowledge (not knowledge-how).

and we could call this (although I will not generally do so) an *absolute illumination* conception of knowledge.<sup>15</sup>

<sup>15</sup> Even if a particular epistemologist insists that she subscribes only to most, not all, of this conception of knowledge, that will not undermine the spirit of this book’s argument. For I will be questioning *each* element of this conception. And each element will, whenever present at all within an individual epistemologist’s conception of knowledge, be *centrally* present. Maybe a better name than ‘the standard analytic conception’ would be ‘a paradigmatic analytic

That conception could be distilled into a conjunction of these five theses: *Beliefism*, *Justificationism*, *Absolutism*, *Gettierism*, and *Theoreticalism*. *Beliefism* is Section 1.1's knowledge-as-either-a-belief-or-an-unanalysable-state thesis. *Justificationism* is Section 1.2's knowledge-as-well-supported thesis. *Absolutism* is Section 1.3's thesis of knowledge-absolutism. *Gettierism* is Section 1.4's knowledge-as-not-Gettiered thesis. And *Theoreticalism* is Section 1.5's knowledge-that-as-fundamentally-theoretical-knowledge thesis. Already, we have noticed other theses clustering around these ones. But the conjunction of these five adequately conveys the absolute illumination conception.

## 1.7 *Prima Facie* Core Problems

The standard analytic conception of knowledge might, or might not, have fully precise boundaries.<sup>16</sup> I take no stand on that. Nevertheless, even if its boundaries can be somewhat vague, it is intended to have some sharply distinguishing features — some vital marks, some core components. Already, we have identified a few of these.

*Must* they be part of a correct conception of knowledge, though? The rest of this chapter will introduce three *prima facie* challenges to the standard analytic conception of knowledge. I call these *prima facie core* problems, because if they are real problems they are not conceptually peripheral ones. They concern the heart of that conception. Chapter 2, Chapter 3, and Chapter 4 will investigate each in turn, in some detail.<sup>17</sup>

---

conception'. Even so, I trust, no serious assessment of the book's merits will depend upon this choice of name.

<sup>16</sup> If it does, it is what Lynch (1998: 57) calls 'crystalline' and what Battaly (2001: 105) calls 'Fregean'. Such concepts, says Lynch,

*are like crystal*: rigid, pure, and transparent, with sharp edges and definite borders.

Battaly describes such a concept as being 'maximally thick':

A concept is maximally thick when all of the necessary and sufficient conditions for its application are fixed — when all of its boundaries are precise.

So, sharp boundaries are needed, at all points, if a concept is to be Fregean or crystalline.

<sup>17</sup> Chapter 5 will then address a consequent challenge: *If* these are genuine core problems, collectively they constitute a problem not only for the content of the standard analytic conception of knowledge, but also possibly for the view that knowledge can be analysed at all. In recent years, this concern about the viability of any conceptual analysis of knowledge has been urged by Williamson (2000: ch. 1). A question that emerges from his work, therefore, is this: Can we understand knowledge non-reductively, in terms just of knowledge? Chapter 5 will suggest a way of doing so.

### 1.7.1 The justificationism problem<sup>18</sup>

Section 1.2 alerted us to the widespread epistemological acceptance of justificationism. This is the knowledge-as-well-supported thesis, the usual (partial) precisification of which claims that knowledge must include some form of justification — not necessarily an evidential form, but *some* form of *something* epistemising. It is no less clear, though, that there is almost equally widespread reticence as to quite how *much* justificatory support a true belief needs if it is to be knowledge. Voluminous philosophical discussion exists also, concerning what kind of justification is required. (Witness the energetic debates on the respective merits of such ideas as evidentialism, reliabilism, defeasibility, contextualism, coherentism, and foundationalism.) But how *strong*, in particular, must knowledge's justificatory component be? On this, there is almost silence.

At any rate, that is true of proponents of fallibilist theories of knowledge.<sup>19</sup> Infallibilists about knowledge accept that a true belief is not knowledge unless the believer has justificatory support for it (even if in a broad sense of 'justificatory') that leaves no rational possibility of its being false. Justificatory support needs to be perfect in that respect (they say); otherwise, there is not really knowledge present.

Yet fallibilists probably constitute the overwhelming majority of epistemologists,<sup>20</sup> and they spurn such justificatory elitism. They assure us that a true belief can be knowledge even if its justificatory component provides merely fallible support for the belief's being true. And this is a heartening idea if we regard it as able to be part of a coherent, indeed correct, conception of knowledge according to which much knowledge is widely available.

Nonetheless, that optimistic fallibilist thesis could well be mere wishful thinking if fallibilists are unable to say *how* fallibly a true belief can be justified without falling short of being justified enough to be knowledge.

<sup>18</sup> Part of Section 1.7.1 is from Stephen Hetherington, 'Knowledge's Boundary Problem', *Synthese* 150: 41–56. Springer, 2006. Reprinted by permission of Springer.

<sup>19</sup> Or almost so. The problem is rightly regarded by Bonjour (2002: 43, 46, 48–9; 2003: 21–3) as epistemologically fundamental. But his response is to adopt an *infallibilism* about knowledge.

<sup>20</sup> There are not many sceptics among epistemologists. And infallibilists tend to be sceptics. After all, we have few, if any, infallibly justified beliefs. So, there is little, if any, knowledge — provided that knowledge must include infallible justificatory support. (For more on fallibilism, see Hetherington 2005a; Vahid 2008; Fantl and McGrath 2009.) Still, it is not always clear whether a particular epistemologist is a fallibilist. Some epistemologists, it seems, regard the idea of fallible knowledge as incoherent or empty (e.g., Lewis 1996; Hendricks 2006: 9) — even while extensionally restricting, in each given case, what the notion of infallibility is encompassing.

How fallible is too fallible, in that respect?<sup>21</sup> Unfortunately, all that is usually suggested is that a true belief's being knowledge involves its being well justified (or other words are used, to similar effect). And what — either precisely or even helpfully-but-imprecisely — does that mean? We are yet to be told.

Accordingly, epistemologists in general face a conceptual challenge of either removing or disarming that vagueness within any fallibilist conception of knowledge. Most epistemologists need to show why that vagueness does not undermine all putative fallibilist theories of knowledge. Elsewhere (2006a), I have called this challenge knowledge's boundary problem. Here, I call it knowledge's *justificationism problem*. It is the epistemological problem of knowing how much fallibility is allowable within knowledge's required justificatory component. Traditional fallibilist epistemology is confident that there must be some lower bound on the amount and quality — for short: the strength — of justification sufficient to distinguish a belief's being true from its being knowledge. Yet fallibilist epistemologists are far from agreeing on where that boundary lies. And their problem is not simply one of achieving consensus. The situation is not one where we find a plethora of individual epistemologists vigorously defending their respective delineations of knowledge's justificatory boundary — before disagreeing with each other over its precise location, thus failing to reach a consensus. Instead, almost all epistemologists are simply silent on these details. It is as if they do not even realise that knowledge has a justificationism problem like this.<sup>22</sup>

Yet their apparent inability to locate the justificatory boundary between knowledge and non-knowledge is a ground for doubt on their part as to the boundary's very existence. More pointedly, it is a ground for doubt with the following features. First, it is not a remote ground for doubt. It is a realistic doubt, with epistemologists (both collectively and individually) actually — not just possibly — being unable to say non-arbitrarily where knowledge's justificatory boundary is located. Second, it is a central doubt, not a peripheral one, pertaining as it does to one of knowledge's supposedly defining

<sup>21</sup> And would any given fallibilist epistemologist, if confronted with the actual justificatory boundary between knowledge and non-knowledge, know that it — *as against a minimally different putative justificatory boundary* — is the boundary? It is hard to imagine how a minimally different putative justificatory boundary would ever justifiably seem to the epistemologist to be less likely to be the actual justificatory boundary.

<sup>22</sup> I mentioned BonJour earlier. Fantl (2003: 559) is another who does realise this. His own preference is for a fallibilist infinitism. That option will be discussed briefly in Section 5.6. Fantl notes also that Rudner (1953) and Owens (2000: 25–6) appreciate how potentially significant is the apparent unavailability of a non-arbitrary yet fallibilist justificatory boundary, not only for knowledge-versus-non-knowledge but for acceptance-versus-non-acceptance (Rudner) or belief-versus-non-belief (Owens).

characteristics. Third, it remains uneliminated. Collectively, it has not yet been eliminated by epistemologists. Perhaps no individual epistemologist, too, has eliminated it.

But when a fallibilist epistemologist encounters a realistic, central, and currently uneliminated ground for doubt, how should he react? He must take seriously the possibility of there being a correlative lack of knowledge on his part. And if he cannot eliminate that ground for doubt, he should infer that there is a correlative lack of knowledge: in general, a fallibilist should infer — from his noticing a presently undefeated, realistic, and central ground for doubt as to  $p$  — that there is a lack of knowledge that  $p$ .<sup>23</sup> This inferential reaction can be reversed, of course, if the ground for doubt is subsequently defeated. In the meantime, though, knowledge departs. So, there is a special reason for fallibilists to be perturbed at the existence of a realistic, central, and not yet defeated ground for doubt about the location of knowledge's justificatory boundary. And most epistemologists are fallibilists about knowledge. Far from unworriedly presuming that knowledge has a justificatory boundary, therefore, perhaps most epistemologists should be inferring that *they do not know* there to be a justificatory boundary, even an imprecisely described one,<sup>24</sup> between knowledge and non-knowledge — a boundary constituted by the presence or the absence of some strength of some sort of justification.<sup>25</sup>

Significantly (and as we have seen just now), this result follows from those fallibilists' own ways of conceiving of knowledge. It is not a result being imposed on them by an infallibilism about knowledge. It does not reflect simply their not satisfying an infallibilist conception of what is required to know that knowledge, as part of its core, has a justificatory boundary.

<sup>23</sup> The same is true, *mutatis mutandis*, of infallibilists. In general, an infallibilist should infer, from there being any undefeated ground at all for doubt as to  $p$ , that there is a lack of knowledge that  $p$ . Arguably, fallibilists need to give way only to undefeated grounds for doubt that are realistic and central; definitely, infallibilists do not require the grounds for doubt to be realistic or central.

<sup>24</sup> The point would not be that only a very *precisely* described justificatory boundary eludes us. Even a definitively described imprecise-but-not-too-much-so justificatory boundary would do so.

<sup>25</sup> 'Is that too strong a claim, given the evidence? This analogy could suggest so. We cannot, in general, determine exactly what level of wealth demarcates rich people from non-rich ones. Still, this should not make us doubt that there *is* a real boundary between the rich and the non-rich.' Yet I am not claiming there to be no difference between being justified and not being justified. My criterion is that possibly *this* difference is not what distinguishes knowing from not knowing. Maybe a better analogy is as follows. Possibly, what distinguishes the happy (knowing) from the unhappy (not knowing) is not money (justification) — because, for a start, we cannot ever find a monetary (justificatory) boundary that adequately demarcates that difference. The correlative *question* is whether knowledge (happiness) is such that justification (money) is part of it at all. (And I am about to refine that question.)

Rather, they are failing a fallibilist standard: they are yet to eliminate a realistic and central ground for doubt as to whether knowledge has a core justificatory boundary. Accordingly, fallibilists should be denying themselves the knowledge of there being such a required justificatory boundary for knowledge.<sup>26</sup>

We thus have a *prima facie* challenge to the idea that knowledge has to include justification. Yet what would knowledge be like, if it was not required to include justification? Can we make conceptual sense of that suggestion? Might the suggestion even be correct? Chapter 4 will show how it might well be. But to defend that possibility is to defuse a key component of the standard analytic conception of knowledge. Epistemologists tend to argue for that component — the knowledge-as-well-supported thesis — by (1) describing instances of true but unjustified beliefs before (2) telling us how obvious or intuitive it is that no such belief is knowledge. Chapter 4 will challenge both the apparent pre-theoretical innocence and the epistemic security of that approach, in part by uncovering something quite philosophically substantial that we have been taking for granted — but that we might well not know to be true — when we insist upon knowledge's including some sort of good epistemic support. As we will find, the usual epistemological commitment to the knowledge-as-well-supported thesis is not as philosophically unquestionable as we may wish it to be. Fortunately, we will also find that this does not entail our having to become sceptics about there being any knowledge. However, we will need to modify a central part of our standard analytic conception of what knowledge must be like.

<sup>26</sup> That is unfortunate enough; does fallibilism's plight then worsen? A commitment to knowledge's having a justificatory boundary is a vital part of all traditional fallibilist conceptions of what knowledge is. That commitment is the firm, maybe unshakeable, acceptance of justificationism — the knowledge-as-well-supported thesis. Consequently, those conceptions should at least confront the thought that anyone who lacks knowledge of knowledge's having a justificatory boundary at its core might well also lack knowledge of knowledge's even existing in the first place. Here we may reach for Quine's influential and insightful dictum, 'No entity without identity' (e.g., 1969: 23). Does a lack of knowledge as to where to locate knowledge's justificatory boundary imply a lack of knowledge as to knowledge's identity conditions? And if we do not know knowledge's identity conditions, can we know that knowledge even exists? This dramatic thought is inessential to my argument; I mention it for the following reasons. *Prima facie*, if there is knowledge (and we are confident of this), we should at least be less confident of the knowledge's existing in part *because* of a justificatory boundary between knowing and not knowing. And then, *if* we remain committed to knowledge's requiring a justificatory boundary between it and not knowing, we could well begin to wonder whether we really understand what knowledge is. At which point, we could well begin to wonder whether what we have taken to be knowledge is really knowledge. At which point, also (and more dramatically), we could well begin to wonder whether there really is knowledge.



### 1.7.2 The Gettierism problem<sup>27</sup>

No contemporary epistemological discussion of the nature of fallible knowledge may responsibly ignore the project of distinguishing (1) having knowledge that  $p$ , from (2) failing, *due to being in a Gettier situation*, to have knowledge that  $p$ . As Section 1.4 implied, some will think of this project as a special case of the project mentioned in Section 1.2, of determining what kind and amount of justification is required within knowledge. Hence, equally, some should wonder whether Section 1.7.1's justificationism problem has, as a special instance, a *Gettierism problem*. We may also wonder as to this on independent grounds. Is the standard analytic commitment to Gettierism true? Or is there a problem of principle in maintaining that commitment within a conception of knowledge?

Certainly, there has been a conspicuous lack of agreement among epistemologists as to how to specify the difference between a fallibly justified true belief's being Gettiered (and thereby not being knowledge) and its not being Gettiered (and thereby being knowledge). Yet their collective confidence that there is some such difference — what could be called a Gettier-boundary between fallibly knowing that  $p$  and fallibly not knowing that  $p$  — has remained undaunted. Perhaps it is now appropriate to question that general confidence. Might it have settled, albeit imperceptibly, into a form of dogmatism? Maybe we should consider, instead, the possibility that we will *never* be able to find a non-arbitrary specification of the Gettier-boundary (and thereby of Gettierism's requirement that there *be* some such boundary). What would follow from that impossibility? Might we have been misinterpreting, all along, what it is within Gettier situations to which we have been responding 'intuitively'? Might we have taken Gettier situations to be revealing something that, as it happens, they do not?

I will illustrate that general question via a few examples, beginning (purely for specificity) with defeasibility theories of knowledge.

The basic idea within such theories is that a justified true belief is knowledge only when its justification component is undefeated. That idea is meant to be interpreted in a fallibilist way (even though the absence of defeat is sometimes called complete justification: e.g., Lehrer and Paxson 1969). The fallibilism enters the story through the latter's description of the phenomenon of a defeater: only relevant or salient circumstances, say, are potential defeaters. Then that fallibilism is combined with a completeness requirement: all such circumstances need to be defeated.<sup>28</sup>

<sup>27</sup> Or, as Lycan (2006) calls it, 'the Gettier Problem problem'.

<sup>28</sup> For a recent, more detailed, formulation of a defeasibility condition, see Bergmann (2006: ch. 6).

Now, that might sound like a sufficiently-well-described justificatory state. Nonetheless, there is inherent vagueness in the key idea behind defeasibility analyses. Defeaters are circumstances that defeat by weakening justification: as more and stronger defeaters are being overlooked by a particular body of evidence, that evidence is correlatively weakened.<sup>29</sup> And how weak — either exactly *or* inexactly — can the justification for a belief become, as a result of some circumstance's presence, before it is too weak to sustain the belief's being knowledge? This amounts to the question, raised by William Lycan (1977), for instance, of how much — and which aspects — of one's environment need to be noticed by one's evidence, if that evidence is to be undefeated, thereby being justification that makes one's belief knowledge (all else being equal). To say the least, it is not always clear, even roughly so, where to draw the line between aspects of the environment that do — and those that do not — need to be countered by one's evidence. What should we expect of people in this respect? No non-arbitrary (and fallibilist)<sup>30</sup> answer suggests itself.

Accordingly, it is unsurprising that, in evaluating Fred Dretske's early instantiation (1971) of the general category of 'tracking theories of knowledge', Fred Adams (2005: 21; my emphasis) has this observation:

True, there is some vagueness over how close the relevant alternatives (Trudy) can be before Tom loses knowledge that Judy is in front of him. *This is a general difficulty about metrics for nearness of possible worlds, when testing counterfactuals.*

Indeed it is. Like other epistemologists, Adams is not attending to the possibility of facing a Gettierism problem as such — that is, a problem due to his requiring, from the outset, that any instance of knowing not be Gettiered. Nonetheless, he notices at least the data — the vagueness, the inherent difficulties, in specifying what it is to not be Gettiered — that constitute a *prima facie* case for the existence of the Gettierism problem. How relevant, how salient, must alternative possible circumstances be, if they are to assist in constituting the presence or the absence of a particular piece of knowledge? Epistemologists are aware of the relevance, the salience, of that question. Yet is

<sup>29</sup> This is so, even when the defeaters clash directly with one's belief. And it is so, regardless of the believer's not realising that the evidence is thereby weakened.

<sup>30</sup> In contrast, infallibilism could be thought to be describing a non-arbitrary standard — because there is only one possible justificatory strength it accepts as being sufficient (other things being equal) to satisfy all justificatory demands within an instance of knowledge. (Fallibilism has to choose between many possible such justificatory strengths.) But infallibilism, considered in opposition to fallibilism's plethora of options, could be an arbitrary strength to favour in the first place. In any case, infallibilism is not the strength being tested within Gettier cases. (Gettier himself was explicit about this (1963: 121). His cases pertained only to justification providing fallible support.)

there no non-arbitrary answer to it? Epistemologists have not fully confronted the possibility that there might not be any such answer — and that this might reflect knowledge's having no Gettier-boundary in the first place.

However, that possibility should be taken seriously by epistemologists, because the problem adverted to by Adams applies, *mutatis mutandis*, to all proposals for solving the Gettier problem. To take another example: If we are told that a necessary condition of a belief's being knowledge is its being formed and maintained in some causally apt way (e.g., Goldman 1967), we should then ask this sort of question: 'How apt, how non-deviant, how well-behaved — even imprecisely so — must that causal nexus be?' And we should not expect a definitive answer to our question. No one has yet offered one; nor do I assume that this will occur. There might well be no unequivocally delineated dividing line, even one drawn only roughly, between (1) accurate descriptions of how a belief may be formed so as to be knowledge (in part, by avoiding the 'Gettier trap') and (2) accurate descriptions of how it may not. Must there be some such causally specifiable Gettier-boundary? Maybe not.

So, the possible conceptual danger appears to be quite general, even one of principle, afflicting all extant and all potential Gettier-sensitive epistemological descriptions of knowledge. Perhaps no non-arbitrary answer (even only a roughly correct one) to the question of where the Gettier-boundary is located could ever impress itself upon epistemological minds. At the very least, there is already some substantial inductive evidence for thinking that no such answer will ever present itself to epistemologists. They are confident, as a group, that knowledge is absent from Gettier cases. But epistemologists are also unconfident, as a group, that they really understand why knowledge is absent from Gettier cases. Hence, not only are they yet to agree on where to locate the Gettier-boundary; it is easily conceivable that they will never agree on it. Nevertheless, they continue to insist that there is some such boundary. They continue to insist that some such boundary is required by the core of their conception of knowledge. The following question must therefore be asked: At what stage, if ever, in their continuing struggles to specify that Gettier-boundary should epistemologists begin wondering whether their insistence on its existence in the first place is mistaken?<sup>31</sup>

That confronting option would constitute a possibility of epistemologists in general having miscategorised Gettier cases in an especially fundamental way. Still, why could that not occur (even if it is not *probably* what has

<sup>31</sup> 'It feels so intuitive, though, to say that the core of our conception of knowing reveals the existence of a Gettier-boundary for knowing.' Perhaps so, at least for now and for epistemologists. But see Section 3.14 on epistemological intuitions and Gettier. (In Chapter 3, my interpretation of Gettier's challenge will be based on a balancing of theoretical considerations, not on intuitions.)

happened)?<sup>32</sup> What are often termed ‘Gettier intuitions’ are the standard, supposedly intuitive, evaluations by epistemologists of Gettier cases — those evaluations that categorise the cases as ones from which knowledge is missing. Yet such ‘intuitions’ are responses to what are actually quite complex situations: Gettier cases contain several epistemic variables. Moreover, although those ‘intuitions’ have been formulated in terms that strike epistemologists as the best of the available conceptual alternatives, maybe we have collectively conceived of, let alone consulted, too few such alternatives when trying to articulate our shared sense of there being something amiss in Gettier cases. That thought will be tested in Chapter 3. It is a potentially invigorating thought, too. For we may interpret it optimistically, as encouraging us to search for new conceptual alternatives with which to describe what goes awry within Gettier cases. Perhaps we need to expand the conceptual repertoire with which we respond to such cases.

If we can do so, will that locate the Gettier-boundary for us, allowing the retention of Gettierism at the core of our conception of knowledge? I doubt that it will; the expansion we must contemplate is more radical still, I fear. The quest to distinguish fallible knowledge from fallible non-knowledge, in part by distinguishing non-Gettiered fallibly justified true beliefs from Gettiered ones, has so far been unsatisfied by epistemologists in general. And, I urged just now, this lack of collective success could well be a problem of principle: possibly it is inescapable for us, no matter what concepts we reach for in our efforts to describe the Gettier-boundary. As we ponder possible explanations for that unsatisfying history, we should consider this one: Perhaps there is *no* Gettier-boundary for cases of knowledge. Considered in relation to post-1963 epistemology as a whole, this is a radical hypothesis; it should not therefore be rejected out of hand, though. Nor, equally, need it be accepted immediately. As I said, Chapter 3 will test it. We will need to consider carefully afresh what knowledge is like, if it is not necessarily to be understood (even in part) through a need to isolate a Gettier-boundary for it.

### 1.7.3 The theoreticalism problem

The justificationism problem and the Gettierism problem arise for propositional or factual knowledge. They are potential problems for our being able to understand — while ever we retain justificationism and Gettierism as

<sup>32</sup> ‘But presumably the fact that epistemologists cannot agree on why there is a lack of knowledge within Gettier cases does not entail that there is no such lack within those cases.’ I agree. Right now, though, I am motivating only a *prima facie* possibility — of epistemologists having *misjudged* there being a lack of knowledge within Gettier cases, with one sign of this being their continuing inability to concur as to why knowledge is absent from the cases. At *some* stage along the way, we should begin to wonder whether that sort of misjudgement could be what actually explains that absence of concurrence.

standard conceptual commitments regarding — what it is to know fallibly that *p*. Now we should notice another possible core problem — a *theoreticism problem*.

This one concerns the standard epistemological emphasis upon trying first and foremost to understand propositional knowledge, as against other sorts of knowledge. In particular, Section 1.5's knowledge-that-as-fundamentally-theoretical-knowledge thesis is now to be considered — with knowledge being assumed to take a theoretical form, not a practical form. Do epistemologists know that knowledge-that is definitively theoretical, not practical? The worry behind this question is that, in the absence of real knowledge of the claimed categorial difference between propositional knowledge and knowledge-how, the usual epistemological insistence upon there being such a categorial boundary could well be leading us fundamentally astray in our attempts to ascertain even propositional knowledge's key constitutive and distinguishing characteristics. Epistemologists would not know what propositional knowledge is, to the extent that they would not know the respects, if indeed there are any, in which it is not something else — in this case, knowledge-how.

And is there any *prima facie* ground for dissolving that standard theoreticist assumption of the existence of that putative categorial boundary — marked by the theoreticality of knowledge-that, notably distinguishing such knowledge from knowledge-how (practical knowledge)? Chapter 2 answers this question. As will be acknowledged there in more detail, other philosophers have also questioned the distinction. But in a crucial respect they have done this less fully than might be needed. They have retained the traditional assumption of the conceptual priority of propositional knowledge, by seeking to analyse knowledge-how in terms of knowledge-that. Again, Chapter 2 will question that assumption.

Why should we even contemplate questioning it, though? Here are two reasons.

- (1) There is no clear consensus that any existing attempts to analyse knowledge-how in terms of knowledge-that have succeeded. Correlatively, there is no clear consensus that philosophers have shown the conceptual priority of propositional knowledge over knowledge-how.
- (2) Consider the contrary possible priority of knowledge-how over knowledge-that — and a corresponding putative reduction, of knowledge-that to knowledge-how. Such a reduction, if it was to succeed, would displace propositional knowledge from its traditionally assigned pre-eminence within the ontology of categories of knowledge. And *might* that suggested reduction succeed? Is the possibility of this contrary sort of reduction *prima facie* worthy of investigation? I believe so. The following reasoning supports that *prima facie* case (and Chapter 2 will expand upon these remarks).

## 22 *The Standard Analytic Conception of Knowledge*

First, it is clear that, in having some knowledge that  $p$ , one thereby has various cognitive and/or practical skills — various kinds of knowledge-how. We may formulate that general observation slightly more specifically, in these terms:

- K In having the knowledge that  $p$ , one is thereby — to some more or less marked extent — able to appreciate, report on, reason about, and/or generally react to, more or fewer  $p$ -related aspects of the world.

There is vagueness aplenty in K's details, a fact to which I return in a moment. Before then, the following question is immediately pressing. What (if any) constitutive relationships obtain between (1) the knowledge that  $p$  mentioned in K and (2) those abilities also mentioned in K? I used, within K, the phrase 'in having' and the term 'thereby'; should we interpret these in a constitutively strong way? In particular, could it be that one's having these abilities is at least *part of* having the knowledge that  $p$ ? K says that, in having the knowledge, one 'thereby' possesses those abilities. Could this be because the knowledge actually includes the abilities? Is that constitutively possible?

Epistemologists do not talk in that way about what it is to know that  $p$ . Instead, their usual locutions treat each instance of propositional knowledge as being a state *from which* those skills may flow. One would have the knowledge that  $p$  by (1) having various pertinent properties and thus by (2) some apposite state of affairs (the knowledge-state) obtaining — with (1) and (2) being so in ways that would not depend upon the prior or simultaneous presence of those other abilities mentioned in K. These abilities would be present, if at all, only as a *conceptually separable consequence* of the knowledge's presence. In other words, as a matter of conceptual priority the knowing-that would be independent of the associated cognitive abilities — with these, in themselves, being instances of cognitive knowledge-how.<sup>33</sup> The knowledge-that would be constituted prior to the associated instances of knowledge-how.

Second, however, there is a problem, maybe one of principle, in describing (even roughly) where the constitutive boundary is to be located between the knowledge-that and those associated abilities, these cases of knowledge-how. Where does the former end, with the latter then able to flow from it, expressing or manifesting it? There is the realistic prospect of ineliminable arbitrariness lurking within all suggestions we may make as to how to understand any insistence upon the knowledge that  $p$ 's being something distinct from, let alone conceptually prior to, the network of associated abilities, those forms of knowledge-how. Consider the following sample questions. Are any of those abilities essential to a given instance of having the knowledge that  $p$ ?

<sup>33</sup> Let us assume so for now. Chapter 2 will discuss this characterisation in some detail.

Are only some of them essential? (If so, which ones? How many would be required? How strongly must any given ones be present?) Maybe no particular ability is essential to each instance of knowledge that  $p$ ; but perhaps, for each such instance, at least one of those abilities is essential. Is this possible? If it is, how are we to understand more fully that use of ‘at least’?<sup>34</sup>

We need not already have discarded the standard analytic conception of knowledge, if we are to take those questions seriously. The data from which they naturally emerge are elements of that conception (no matter that the latter’s proponents have not focused upon these questions). For example, even to describe knowledge as, in part, a belief (as Section 1.1’s traditional knowledge-as-belief thesis does) is to open the door to these questions.<sup>35</sup> We might well say that to believe that  $p$  is to be disposed to react, in apt ways and to apt extents, to some smaller or larger range of  $p$ -related propositions, questions, or aspects of the world (where these will bear more or less clearly upon  $p$ ).<sup>36</sup> These dispositions are cognitive abilities, albeit of a depth, strength, and truth-conduciveness proportional just to belief (rather than something epistemically stronger). They include semantic recognition, perhaps some sorting of evidence, maybe other introspective capacities. And they are  $p$ -related cognitive abilities, to the extent that they link the believer to  $p$  as such. This could include abilities to respond to questions directly about  $p$ , but also to ones bearing only indirectly upon  $p$ .

However, now this question arises (similar in form to that with which I responded to K):

<sup>34</sup> Consider what Ryle (to whom philosophers owe their greatest insights into the distinction between knowledge-that and knowledge-how) says about this. He aims ‘to prove that knowing-that presupposes knowing-how’ (1971: 224). So he argues, first, that one only ever knows a truth if one has *already* performed actions (e.g., discovering, checking, running tests, etc.) that amount to knowledge-how; and, second, that one knows a truth only if one is *subsequently* able to perform actions that amount to knowledge-how (allowing one to ‘intelligently exploit’ the knowledge-that). For the sake of argument, grant Ryle those two claims. How then will knowledge-that presuppose knowledge-how without *being* knowledge-how? Maybe Ryle is right that there is no ghostly mind in the machine. But is ghostly knowledge-that still lingering? As I will soon explain, we should be willing to investigate the possibility of knowledge-that being knowledge-how, when all is said and done. Ryle is walking a thin line, retaining knowledge-that’s distinctness from knowledge-how — even while conceding these close links, which we might suspect are constitutive ones, between the two.

<sup>35</sup> The same is true, *mutatis mutandis*, of the more general knowledge-as-either-a-belief-or-an-unanalysable-state thesis, also in Section 1.1.

<sup>36</sup> In a related spirit, consider Cohen’s (1992: 4) similarly schematic characterisation of belief:

belief that  $p$  is a disposition, when one is attending to issues raised, or items referred to, by the proposition that  $p$ , normally to feel it true that  $p$  and false that *not-p*, whether or not one is willing to act, speak, or reason accordingly.

For an earlier dispositional analysis of belief, see Armstrong (1973: part I). For useful critical commentary, see Stich (1984).

Does a belief that *p* include — or does it only give rise to — some kind of presence of such *p*-related abilities? Different abilities might strike us as being of greater or lesser relevance here, perhaps depending upon different values of '*p*'. Which, if any, are required within a given case? How strongly must they be present?

To the extent that we think at least some of those abilities are required to at least some extent and with at least some force (such as if we require the believer to recognise the pertinence of a wide range of questions about *p*), we also give correlative credence to the thought that they help to constitute the presence of the belief, and thereby the knowledge-that. Yet do all of them provide equal such help? Which abilities — which of these forms of knowledge-how — are especially important? Where should we draw the conceptual boundary here, even roughly, between those abilities that are part of, and those that are only produced by, a given instance of belief (and thereby a given instance of knowledge-that)?

As was true of the justificationism problem, therefore, the theoreticalism problem confronts us with a series — indeed, possibly several intersecting series — of potential cut-off points. In order to appreciate this prospect more generally, we need only to ponder the vast scope that exists for alternative possible precisifications of K, the description I offered of the associated abilities that one would 'thereby' possess in having an instance of knowledge that *p*. Accordingly, here is a slightly compressed and altered version of that description, emphasising the 'pressure points' from which problems may spring, as we seek marks distinguishing knowledge-that from knowledge-how:

K\* In having the knowledge that *p*, one is thereby, to *some more or less* marked extent, able to do *some or all* of those activities mentioned in K (along with *indeterminately many* others that may equally well have been mentioned), all of which involve one's reacting *more or less well or strongly* to *more or fewer* aspects of the world that are *somehow p*-related.

Thus, much scope is allowed by K\* for competing precisifications of it — and thereby for competing boundaries between the knowledge that *p* and these abilities that would amount to associated knowledge-how.<sup>37</sup>

And so arises the theoreticalism problem for knowledge. Does the existence of that scope regarding K\*'s precisification undermine any claims we

<sup>37</sup> Note that here, as ever throughout Section 1.7 when talking of precisification, I mean to be encompassing even *rough* precisifications. (These are not the same as *imprecisions*, after all.) Consequently, I am not requiring an especially high degree of precisification. Even when seeking only an approximate sense of how to delineate these boundaries, it seems as if we might well face the challenges I have described.



might make of knowing where, even roughly, to locate the potentially complex categorial boundary between knowledge that  $p$  and the associated  $p$ -related kinds of knowledge-how gestured at in  $K^*$ ? If so, should we begin taking seriously, at least so as to think of investigating further, the idea that no such boundary exists? Maybe instead  $K^*$  is correct, in its vague way, as a partial description of how knowledge-that is constituted — hence, as a partial description of how knowledge-that is itself a form of knowledge-how. The next chapter will examine that thought, bypassing the usual search for a description of how the theoreticality of knowledge-that differs from the practicality in knowledge-how.