

# CHAPTER 18

## A State of Mind

---

Timothy Williamson

### 1 Factive Attitudes

Knowing is a state of mind. That claim is central to the account of knowledge developed in this book. But what does it mean?

A state of a mind is a mental state of a subject. Paradigmatic mental states include love, hate, pleasure, and pain. Moreover, they include attitudes to propositions: believing that something is so, conceiving that it is so, hoping or fearing that it is so, wondering whether it is so, intending or desiring it to be so. One can also know that something is so. This book concerns such propositional knowledge. If  $p$  is a proposition, we will understand knowing  $p$  not as merely being acquainted with  $p$  but as knowing that something is so, something that is so if and only if  $p$  is true. For example, if  $p$  is the proposition that it is cold, then one is acquainted with  $p$  in merely wondering whether it is cold; to know  $p$  is to know that it is cold. Knowing in that sense is a *factive* attitude; one knows  $p$  only if  $p$  is true, although one can be acquainted with the proposition  $p$  even if it is false. Other factive attitudes include perceiving that something is so, remembering that it is so, and regretting that it is so. If attitudes are relations of subjects to propositions, then the claim is that

knowing itself is a mental relation such that, for every proposition  $p$ , having that relation to  $p$  is a mental state. Thus for some mental state  $S$ , being in  $S$  is necessary and sufficient for knowing  $p$ . We abbreviate that claim by saying that knowing is a mental state.

We may assume initially that knowing  $p$  entails believing  $p$ ; section 5 considers that assumption in more depth. Someone might expect knowing to be a state of mind simply on the grounds that knowing  $p$  involves the paradigmatic mental state of believing  $p$ . If those grounds were adequate, the claim that knowing is a state of mind would be banal. However, those grounds imply only that there is a mental state being in which is *necessary* for knowing  $p$ . By contrast, the claim that knowing is a state of mind is to be understood as the claim that there is a mental state being in which is *necessary and sufficient* for knowing  $p$ . In short, knowing is *merely* a state of mind. This claim may be unexpected. On the standard view, believing is merely a state of mind but knowing is not, because it is factive: truth is a non-mental component of knowing.

Our initial presumption should be that knowing is a mental state. Prior to philosophical theory-building, we learn the concept of the mental by examples. Our paradigms should include propositional attitudes such as believing and desiring, if our conception of the mental is not to be radically impoverished. But factive attitudes have so many similarities to the non-factive attitudes

that we should expect them to constitute mental states too; we expect a concept to apply to whatever sufficiently resembles its paradigms. It would be strange if there were a mental state of fearing but no mental state of regretting, or a mental state of imagining but no mental state of remembering. Indeed, it is not clear that there are any pretheoretic grounds for omitting factive attitudes from the list of *paradigmatic* mental states. That the mental includes knowing and other factive attitudes is built into the natural understanding of the procedure by which the concept of the mental is acquired. Of course, that does not exclude the subsequent discovery of theoretical reasons for drawing the line between the mental and the non-mental somewhere else. But the theory behind those reasons had better be a good one.

This chapter ... eliminates some putative differences between knowing and non-factive attitudes that might be thought to disqualify knowing as a mental state. The supposed disqualifications concern constitutive dependence on the environment, first-person accessibility, and causal efficacy. In each case, the differences dissolve on inspection. Naturally, this form of argument cannot provide conclusive proof. We survey the current candidates and find them wanting. We can still wonder whether our list of potential differences is complete. But without good theoretical reasons to demote knowing from its pretheoretical status as a central case of a mental state, demotion is surrender to mere special pleading. Indeed, conceptions on which knowing is the wrong kind of state to count as mental are objectionable on independent grounds. We can best understand knowing by classifying it with other mental phenomena.

In this chapter, section 2 orients the claim that knowing is a mental state with respect to some traditional issues about scepticism and self-knowledge. Section 3 explains an incompatibility between the view of knowing as a factive mental state and standard analyses of the concept *knows* as a conjunction of the concepts *believes* and *true* (predicated of the proposition) and of other concepts; it blames the analyses. Section 4 presents a modest positive account of the concept *knows*, distinguishes it from analyses of the traditional kind, and indicates the possibility of understanding epistemology in terms of the metaphysics of

states. Section 5 discusses the relation between knowing and believing, and explores some implications for so-called disjunctive accounts of mental states.<sup>1</sup>

## 2 Mental states, First-Person Accessibility, and Scepticism

The conception of knowing as a mental state can look like a confusion between objective and subjective certainty. Someone might even diagnose that conception as Descartes' central mistake. Did he not seek a mental state sufficient for knowing *p*? Was not clearly and distinctly conceiving *p* his candidate? And does not the failure of his epistemological programme manifest the impossibility of a mental state of the required kind?

On the view to be developed here, if Descartes sought a mental state sufficient for knowing, his mistake lay elsewhere: perhaps in the view (if he held it) that one must always be in a position to know what mental state one is in. H. A. Prichard, who also took knowing to be a mental state, held that one is always in a position to know whether one knows or merely believes (Prichard 1950, p. 86). Few would now claim such powers of discrimination. Indeed, one cause of denials that knowing is a mental state may be the assumption that one must always be in a position to know whether one is in a given mental state.

One is surely not always in a position to know whether one knows *p* (for almost any proposition *p*), however alert and conceptually sophisticated one is. The point is most vivid when the subject believes *p* falsely. Consider, for example, the situation of a generally well-informed citizen N.N. who has not yet heard the news from the theatre where Lincoln has just been assassinated. Since Lincoln is dead, he is no longer President, so N.N. no longer knows that Lincoln is President (knowing is factive). However, N.N. is in no position to know that anything is amiss. He continues reasonably to believe that Lincoln is President; moreover, this seems to him to be just another item of general knowledge. N.N. continues reasonably to believe that he knows that Lincoln is President. Although N.N. does not know that Lincoln is President, he is in no position to know that he does not know that Lincoln is President (see also Hintikka 1962, 106 and section 8.2).

The argument as stated assumes that no a priori reasoning demonstrates that it is impossible to have knowledge about the external world, for such reasoning would make it unreasonable for N.N. to believe that he knows that Lincoln is President. Of course, if all knowledge is impossible then, for any proposition  $p$  whatsoever, one does not know  $p$  and is not in a position to know that one fails to know  $p$ ; one is never in a position to know whether one knows  $p$ . A sceptic about the external world who is not a sceptic about everything might attempt to maintain that, for any informative proposition  $p$  about the external world, one is in a position to know that one does not know  $p$ . Let us assume for the time being that such a sceptic is wrong. ...

We can also construct cases in which one knows  $p$  without being in a position to know that one knows  $p$ . They involve more delicate issues. It is enough for present purposes that one can fail to know  $p$  without being in a position to know that one fails to know  $p$ .

Let transparency be the thesis that for every mental state  $S$ , whenever one is suitably alert and conceptually sophisticated, one is in a position to know whether one is in  $S$ . Given transparency, knowing  $p$  is not a mental state, for almost any proposition  $p$ .

Transparency is false, however, and demonstrably so by reference to uncontentionably paradigmatic mental states. For example, one is sometimes in no position to know whether one is in the mental state of hoping  $p$ . I believe that I do not hope for a particular result to a match; I am conscious of nothing but indifference; then my disappointment at one outcome reveals my hope for another. When I had that hope, I was in no position to know that I had it. Indeed, it is hard to find a non-trivial mental state for which transparency holds. It fails for the state of believing  $p$ , for the difference between believing  $p$  and merely fancying  $p$  depends in part on one's dispositions to practical reasoning and action manifested only in counterfactual circumstances, and one is not always in a position to know what those dispositions are. Transparency is even doubtful for the state of being in pain; with too much self-pity one may mistake an itch for a pain, with too little one may mistake a pain for an itch. ... But even if transparency does hold for a few mental states, it clearly fails for others; the premise of the argument

from transparency to the denial that knowing  $p$  is a mental state is false. Given that knowing  $p$  is a mental state, we will not expect knowing whether one is in it to be always easy.

It does not follow that there is no asymmetry at all between knowledge of one's own mental states and knowledge of the mental states of others. Perhaps failures of transparency could not be the normal case, although that claim would require extensive argument. A more plausible claim is that we have some non-observational knowledge of our own mental states and not of the mental states of others. But then the same may be said of knowing: we have some non-observational knowledge of our own knowledge and ignorance and not of the knowledge and ignorance of others. Any genuine requirement of privileged access on mental states is met by the state of knowing  $p$ . Knowing is characteristically open to first-person present-tense access; like other mental states, it is not perfectly open.

Some may object that knowing whether one knows  $p$  requires evaluating reasons for and against  $p$  in a way in which knowing whether one believes  $p$  does not. They distinguish knowing whether one currently believes  $p$  from deciding whether to continue believing  $p$ . Suppose for a moment that they are correct in taking knowing whether one believes  $p$  not to require one to evaluate reasons for and against  $p$ . Still, even on their view there is also the mental state of *rationaly* believing  $p$ , on some appropriate concept of rationality. Knowing whether one *rationaly* believes  $p$  does require one to evaluate reasons for and against  $p$ . Thus the need for such evaluation in order to know whether one knows  $p$  does not show that knowing  $p$  is not a mental state.

Could it be replied that knowing and *rationaly* believing are not mental states in the way that believing is, because "know" and "rational" are normative terms? Belief attributions have a normative element too, for to have any mental attitude to a content one must in some sense grasp that content, and therefore have some minimal ability to deal *rationaly* with it; the reply itself classifies "rational" as a normative term. In any sense in which "know" and "rational" are normative terms, ascriptions of mental states can be normative.

A different objection is that one's belief about whether one knows  $p$  is defeasible by new information in a way in which one's belief about whether one believes  $p$  is not. For example, the

new information might show that  $p$  is false. But is one's belief about whether one believes  $p$  really indefeasible by new information? Someone might believe that he believes that the world will end next year, because he has joined a religious sect in which there is strong pressure to believe that the world will end next year, but his unwillingness to cash in his pension may suggest that he does not really believe that the world will end next year. When he reflects on his unwillingness to cash in his pension, he may come to that conclusion himself. But even if we forget such examples and suppose that one's belief about whether one believes  $p$  is not defeasible by further evidence, we must still acknowledge mental states such as being alert or thinking clearly about a problem. One's belief about whether one is alert or thinking clearly about a problem is defeasible by new information, for example about what drugs had been slipped into one's drink. Thus the defeasibility of beliefs about whether one knows  $p$  does not show that knowing  $p$  is not a mental state.

Once we consider the full variety of acknowledged mental states, it is clear that any general requirements of privileged access on mental states are very mild. Knowing satisfies those mild requirements.

The failure of transparency helps to clarify the relation between the thesis that knowing is a mental state and a traditional pattern of sceptical argument. The sceptic argues that a subject with a true belief could have been in exactly the same mental state (that is, in the same total set of mental states) even if the belief had been false. He concludes that, since the belief fails to constitute knowledge in the latter case, it fails equally to do so in the former. The sceptical argument assumes something like this: if one's mental state is exactly the same in two situations, then one's knowledge is also the same. On the account to be developed here, that assumption is correct, although not quite in the way that the sceptic imagines.

The sceptic supposes that a difference in knowledge would require some *prior* difference in mental state, which the subject could detect. On the present account, a difference in knowledge would *constitute* a difference in mental state. This difference need not be detectable by the subject who lacks knowledge. Thus the sceptic's assumption is correct for reasons that undermine his argument. He claims to have constructed a

case in which the belief is false although the mental state is exactly the same. But the most that he has really shown about the case is that the belief is false and one's situation is not discriminably different. He has not shown that one cannot be in different mental states in indiscriminable situations. Indeed, since we are sometimes in no position to know whether we are in a given mental state, as argued above, surely one can be in different mental states in situations between which one cannot discriminate (see McDowell 1982).

If knowing is a mental state, then the sceptical argument is not compelling. Indeed, such a view of knowledge need only be defensible for the sceptical argument not to be compelling. Thus *one* route into scepticism is blocked. It is not the purpose of this chapter to argue that all are. ...

If someone has already taken the route into scepticism offered by that fallacious argument, before it was blocked, and has become genuinely undecided, at least in principle, as to whether she is in a sceptical scenario, then the blocking of the route now comes too late to rescue her. Nothing said here should convince someone who has given up ordinary beliefs that they did in fact constitute knowledge, for nothing said here should convince her that they are true. The trick is never to give them up. This is the usual case with philosophical treatments of scepticism: they are better at prevention than at cure. If a refutation of scepticism is supposed to reason one out of the hole, then scepticism is irrefutable. The most to be hoped for is something which will prevent the sceptic (who may be oneself) from reasoning one into the hole in the first place.

The purpose of these remarks has been to give a feel for the view that knowing is a state of mind. The content of the view must now be examined more explicitly. The notion of a mental state will not be formally defined, for that would require a formal definition of the mental. Rather, reflection on the intuitive notion of a mental state will help to clarify its workings. Section 4 will provide a less informal account.

### 3 Knowledge and Analysis

To call knowing a mental state is to assimilate it, in a certain respect, to paradigmatic mental states such as believing, desiring, and being in pain. It is

also to contrast it with various non-examples of mental states. Perhaps the most revealing contrast is between knowing and believing truly.

Believing  $p$  truly is not a mental state, at least, not when  $p$  is an ordinary contingent proposition about the external environment. Intuitively, for example, there is no mental state being in which is necessary and sufficient for believing truly that it is raining (that is, for believing while it is raining that it is raining), just as there is no mental state being in which is necessary and sufficient for believing while Rome burns that it is raining. There is a mental state of believing that it is raining, and there is – on the present account – a mental state of knowing that it is raining, but there is no intermediate mental state of believing truly that it is raining. Let  $S_1$  be knowing that it is raining,  $S_2$  be believing truly that it is raining, and  $S_3$  be believing that it is raining. Then, we may assume, necessarily, everything that is in  $S_1$  is in  $S_2$ ; necessarily, everything that is in  $S_2$  is in  $S_3$ . Nevertheless, on the present account, although  $S_1$  and  $S_3$  are mental states,  $S_2$  is not a mental state.

That something sandwiched between two mental states need not itself be a mental state is not as paradoxical as it may sound. Consider an analogy: the notion of a geometrical property. For these purposes, we can understand geometrical properties to be properties possessed by particulars in physical space. Let  $\pi_1$  be the property of being an equilateral triangle,  $\pi_2$  the property of being a triangle whose sides are indiscriminable in length to the naked human eye, and  $\pi_3$  the property of being a triangle. Necessarily, everything that has  $\pi_1$  has  $\pi_2$ , because lines of the same length cannot be discriminated in length; necessarily, everything that has  $\pi_2$  has  $\pi_3$ . Nevertheless, although  $\pi_1$  and  $\pi_3$  are geometrical properties,  $\pi_2$  is not a geometrical property, because it varies with variations in human eyesight. Something sandwiched between two geometrical properties need not itself be a geometrical property. Similarly, there is no structural reason why something sandwiched between two mental states should itself be a mental state.

The point is general. If  $S$  is a mental state and  $C$  a non-mental condition, there need be no mental state  $S^*$  such that, necessarily, one is in  $S^*$  if and only if one is in  $S$  and  $C$  obtains. The non-existence of such an  $S^*$  is quite consistent with the existence of a mental state  $S^{**}$  such that,

necessarily, one is in  $S^{**}$  only if (but not: if) one is in  $S$  and  $C$  is met. A mental state can guarantee that conjunction only by guaranteeing more than that conjunction.

If the denial that believing truly is a mental state does not immediately convince, think of it this way. Even if believing truly is a mental state in some liberal sense of the latter term, there is also a more restrictive but still reasonable sense in which believing truly is not a mental state but the combination of a mental state with a non-mental condition. The present claim is that knowing is a mental state in every reasonable sense of that term: there is no more restrictive but still reasonable sense of “mental” in which knowing can be factored, like believing truly, into a combination of mental states with non-mental conditions. A sense of “mental” is reasonable if it is sufficiently close to an ordinary sense of the word in important respects. Although the present claim is therefore vague, it is at least clear enough to be disputed.

Strictly speaking, we must distinguish a conceptual and a metaphysical contrast. The conceptual contrast is that the concept *knows* is a mental concept while the concept *believes truly* is not a mental concept. The metaphysical contrast is that knowing is a mental state while believing truly is not a mental state.

The concept *mental state* can at least roughly be defined in terms of the concept *mental concept of a state*: a state is mental if and only if there could be a mental concept of that state. This definition does not in principle exclude the possibility of a non-mental concept of a mental state, for different concepts can be of the same state. We may reasonably assume that states  $S_1$  and  $S_2$  are identical if and only if necessarily everything is in  $S_1$  if and only if it is in  $S_2$ . In a given context, distinct concepts may be necessarily coextensive. For example, since gold is necessarily the element with atomic number 79, the state of having a tooth made of gold is the state of having a tooth made of the element with atomic number 79, but the concept *has a tooth made of gold* is not the concept *has a tooth made of the element with atomic number 79*. Similarly, for any mental state  $S$ , the concept *is in S and such that gold is the element with atomic number 79* is necessarily coextensive with the concept *is in S*, so they are both concepts of  $S$ .

Of the conceptual and metaphysical contrasts, neither immediately entails the other. If the concept *knows* is mental while the concept *believes truly* is not, then it follows immediately that knowing is a mental state, but it does not follow immediately that believing truly is not a mental state, for perhaps there could also be a mental concept of the state of believing truly. Thus the conceptual contrast does not immediately entail the metaphysical contrast. If knowing is a mental state and believing truly is not a mental state, then it follows immediately that the concept *believes truly* is not mental, but it does not follow immediately that the concept *knows* is mental, for perhaps there could be a different concept of the state of knowing which was mental. Thus the metaphysical contrast does not immediately entail the conceptual contrast. Nevertheless, it is hard to see why someone should accept one contrast without accepting the other. If the concept *believes truly* is non-mental, its imagined necessary coextensiveness with a mental concept would be a bizarre metaphysical coincidence. If the concept *knows* were a non-mental concept of a mental state, its necessary coextensiveness with a mental concept would be an equally bizarre metaphysical coincidence. In practice, sloppily ignoring the distinction between the metaphysical and conceptual contrasts is unlikely to do very much harm. Nevertheless, it is safer not to ignore the distinction.

The concept *believes truly* is not a mental concept of a state. If the concept *C* is the conjunction of the concepts  $C_1, \dots, C_n$ , then *C* is mental if and only if each  $C_i$  is mental. For example, the conjunctive concept *is sad and such that gold is the element with atomic number 79* is non-mental, simply because it has the non-mental conjunct *is such that gold is the element with atomic number 79*, although it is a concept of the state of sadness. Even a logically redundant non-mental component concept would make *C* a non-mental concept, although it would then be logically equivalent to a mental concept. By contrast, non-mental concepts in the content clause of an attitude ascription do not make the concept expressed non-mental; the concept *believes that there are numbers* can be mental even if the concept *number* is not. At least, all that is so in a reasonable sense of “mental”, which one might express as “purely mental”. Now the concept *believed truly* is the

conjunction of the concepts *believed* and *true*. The conjunct *true* is not mental, for it makes no reference to a subject. Therefore, the concept *believed truly* is non-mental. Similarly, the concept *believes truly* of subjects rather than propositions is non-mental. The metaphysical and conceptual contrasts turn on whether knowing is a mental state, and on whether *knows* is a mental concept.

Just as the concept *believes truly* is non-mental, so for a similar reason is the concept *has a justified true belief*. Indeed, such an argument applies to any of the concepts with which the concept *knows* is equated by conjunctive analyses of the standard kind. The argument can be generalized to analyses formed using logical connectives other than conjunction. It would not apply if those simpler concepts were all mental, but analyses of the concept *knows* of the standard kind always involve irredundant non-mental constituents, in particular the concept *true*. Consequently, the analysing concept is non-mental: that is, not purely mental. Given that the concept *knows* is mental, every analysis of it of the standard kind is therefore incorrect as a claim of concept identity, for the analysing concept is distinct from the concept to be analysed.

If a non-mental concept were necessarily coextensive with the mental concept *knows*, they would be concepts of the same mental state. The present account does not strictly entail that no analysis of the traditional kind provides correct necessary and sufficient conditions for knowing. But once we accept that the concept *knows* is not a complex concept of the kind traditionally envisaged, what reason have we to expect any such complex concept even to provide necessary and sufficient conditions for knowing?

Experience confirms inductively what the present account implies, that no analysis of the concept *knows* of the standard kind is correct. Indeed, the candidate concepts turn out to be not merely distinct from, but not even necessarily coextensive with, the target concept. Since Gettier refuted the traditional analysis of *knows* as *has a justified true belief* in 1963, a succession of increasingly complex analyses have been overturned by increasingly complex counterexamples, which is just what the present view would have led one to expect.<sup>2</sup>

Even if some sufficiently complex analysis never succumbed to counterexamples, that would

not entail the identity of the analysing concept with the concept *knows*. Indeed, the equation of the concepts might well lead to more puzzlement rather than less. For knowing matters; the difference between knowing and not knowing is very important to us. Even unsophisticated curiosity is a desire to *know*. This importance would be hard to understand if the concept *knows* were the more or less ad hoc sprawl that analyses have had to become; why should we care so much about *that*?<sup>3</sup>

On quite general grounds, one would not expect the concept *knows* to have a non-trivial analysis in somehow more basic terms. Not all concepts have such analyses, on pain of infinite regress; the history of analytic philosophy suggests that those of most philosophical interest do not. “Bachelor” is a peculiarity, not a prototype. Attempts to analyse the concepts *means* and *causes*, for example, have been no more successful than attempts to analyse the concept *knows*, succumbing to the same pattern of counterexamples and epicycles. The analysing concept does not merely fail to be the same as the concept to be analysed; it fails even to provide a necessary and sufficient condition for the latter. The pursuit of analyses is a degenerating research programme.<sup>4</sup>

We can easily describe simple languages in which no necessary and sufficient condition for knowing can be expressed without circularity. Many fragments of English have that property. Why should we expect English itself to be different? Once “know” and cognate terms have been removed, what remains of our lexicon may be too impoverished to frame necessary and sufficient conditions for knowing.

The programme of analysis had its origin in great philosophical visions. Consider, for example, Russell’s Principle of Acquaintance: “*Every proposition which we can understand must be composed wholly of constituents with which we are acquainted*” (Russell 1910–11, at Salmon and Soames 1988, p. 23). Russell calls the principle “the fundamental epistemological principle in the analysis of propositions containing descriptions”. There may well be a reading on which it is correct. However, when the principle is combined with Russell’s extremely intimate conception of acquaintance, it forces analysis to go deeper than the surface constituents of the evidently intelligible propositions of science and common sense, for our acquaintance with those surface constituents is not perfectly

intimate.<sup>5</sup> In such a context, the programme of analysis has a philosophical point. Now the philosophical visions which gave it a point are no longer serious options. Yet philosophers continued to pursue the programme long after the original motivation had gone. Correct deep analyses would doubtless still be interesting if they existed; what has gone is the reason to believe that they do exist.

While the general point is conceded, it might nevertheless be claimed that we have special reason to expect an analysis of *knows*. For we already have the necessary condition that what is known be true, and perhaps also believed; we might expect to reach a necessary and sufficient condition by adding whatever knowing has which believing truly may lack. But that expectation is based on a fallacy. If G is necessary for F, there need be no further condition H, specifiable independently of F, such that the conjunction of G and H is necessary and sufficient for F. Being coloured, for example, is necessary for being red, but if one seeks a further condition whose conjunction with being coloured is necessary and sufficient for being red, one finds only conditions specified in terms of “red”: being red; being red if coloured.

There are other examples of the same phenomenon. Although *x* is a parent of *y* only if *x* is an ancestor of *y*, it does not follow that we implicitly conceptualize parenthood as the conjunction of ancestry with whatever must be added to ancestry to yield parenthood, or even that ancestry is conceptually prior to parenthood. Rather, *x* is an ancestor of *y* if and only if a chain of parenthood runs from *x* to *y* (more formally: if and only if *x* belongs to every class containing all parents of *y* and all parents of its members). Thus parents of *y* are automatically ancestors of *y*. If anything, parenthood is conceptually prior to ancestry; we use the necessary and sufficient condition for ancestry in terms of parenthood to explain why ancestry is necessary for parenthood.<sup>6</sup> Again, *x* is identical with *y* only if *x* weighs no more than *y*, but it does not follow that the concept *is identical with* is the conjunction of *weighs no more than* with whatever must be added to it to yield the former concept, or even that *weighs no more than* is prior to *is identical with*. In this case we explain the entailment by Leibniz’s Law: if *x* is identical with *y*, whatever holds of *x* holds of *y* too, so since *x* weighs

no more than  $x$ ,  $x$  weighs no more than  $y$ . We grasp Leibniz's Law without considering all its instances. In principle one could grasp it before having acquired any concept of weight. Necessary conditions need not be conjuncts of necessary and sufficient conditions in any non-trivial sense.

More generally, the existence of conceptual connections is a bad reason to postulate an analysis of a concept to explain them. For example, the axiom of extensionality says that sets with the same members are identical; it has as good a claim to conceptual truth as the proposition that knowledge entails belief. Nevertheless, the axiom is not explained by an analysis of the concept *set*, if an analysis provides a non-circular statement of necessary and sufficient conditions.

The working hypothesis should be that the concept *knows* cannot be analysed into more basic concepts.<sup>7</sup> But to say that is not to say that no reflective understanding of it is possible.

#### 4 Knowing as the Most General Factive Mental State

Knowing does not factorize as standard analyses require. Nevertheless, a modest positive account of the concept can be given, one that is not an analysis of it in the traditional sense. The one sketched below will appear thin by comparison with standard analyses. That may not be a vice. Indeed, its thinness will clarify the importance of the concept as more complex accounts do not.

The main idea is simple. A propositional attitude is factive if and only if, necessarily, one has it only to truths. Examples include the attitudes of seeing, knowing, and remembering. Not all factive attitudes constitute states; forgetting is a process. Call those attitudes which do constitute states *stative*. The proposal is that 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. Apparent counterexamples to this conjecture are discussed below. The point of the conjecture is to illuminate the central role of the concept of knowing in our thought. It matters to us because factive stative attitudes matter to us.

To picture the proposal, compare the state of knowing with the property of being coloured, the colour property which something has if it has any

colour property at all. If something is coloured, then it has a more specific colour property; it is red or green or. ... Although that specific colour may happen to lack a name in our language, we could always introduce such a name, perhaps pointing to the thing as a paradigm. We may say that being coloured is being red or green or ... if the list is understood as open-ended, and the concept *is coloured* is not identified with the disjunctive concept. One can grasp the concept *is coloured* without grasping the concept *is green*, therefore without grasping the disjunctive concept. Similarly, if one knows that A, then there is a specific way in which one knows; one can see or remember or ... that A. Although that specific way may happen to lack a name in our language, we could always introduce such a name, perhaps pointing to the case as a paradigm. We may say that knowing that A is seeing or remembering or ... that A, if the list is understood as open-ended, and the concept *knows* is not identified with the disjunctive concept. One can grasp the concept *knows* without grasping the concept *sees*, therefore without grasping the disjunctive concept.

We can give substance to the category of factive stative attitudes by describing its realization in a natural language. The characteristic expression of a factive stative attitude in language is a *factive mental state operator* (FMSO). Syntactically, an FMSO  $\Phi$  has the combinatorial properties of a verb. Semantically,  $\Phi$  is an unanalysable expression; that is,  $\Phi$  is not synonymous with any complex expression whose meaning is composed of the meanings of its parts. A fortiori,  $\Phi$  is not itself such an expression.  $\Phi$  also meets three further conditions. For simplicity, they are stated here as conditions on an FMSO in English, although the general category is realized in other languages too. First,  $\Phi$  typically takes as subject a term for something animate and as object a term consisting of "that" followed by a sentence. Second,  $\Phi$  is factive, in the sense that the form of inference from "S  $\Phi$ s that A" to "A" is deductively valid (the scrupulous will read quotation marks as corner quotes where appropriate). Third, "S  $\Phi$ s that A" attributes a propositional attitude to S. On the present view, "know" and "remember" are typical FMSOs. Even with the following glosses, these remarks do not constitute a rigorous definition of "FMSO", but they should make its extension moderately clear.

First, “S  $\Phi$ s that A” is required to have “A” as a deductive consequence, not as a mere cancellable presupposition. There is a use of the verb “guess” on which “S guessed that A” in some sense presupposes “A”. However, this presupposition is cancellable by context, as the logical and linguistic propriety of the following sentences shows:

- (1) I guessed incorrectly that he was guilty.
- (2) I guessed that he was guilty and you guessed that he was innocent.

In contrast, the substitution of “knew” for “guessed” in (1) or (2) yields a contradiction. Incidentally, therefore, the implication from “S does not know that A” to “A” is not like that from “S knows that A” to “A”, for only the former is cancellable. The following sentences are logically and linguistically proper:

- (3) I did not know that he was guilty, for he was innocent.
- (4) I did not know that he was guilty and you did not know that he was innocent.

In contrast, the substitution of “knew” for “did not know” in (3) or (4) yields a contradiction. If  $\Phi$  is an FMSO, the implication from “S  $\Phi$ s that A” to “A” is not cancellable (see Grice 1989, pp. 44–6 and 279–80 for cancellability and the presuppositions of “know” respectively).

Second, FMSOs are stative: they are used to denote states, not processes. This distinction is linguistically marked by the impropriety of progressive tenses. Consider:

- (5) She is proving that there are infinitely many primes.
- (6) The shoes are hurting her.
- \* (7) She is knowing that there are infinitely many primes.
- \* (8) She is believing that there are infinitely many primes.
- \* (9) The shoes are fitting her.

Sentences (7)–(9) are deviant because “know”, “believe”, and “fit” (on the relevant reading), unlike “prove” and “hurt”, are stative. Of course, a verb may have both stative and non-stative readings, as in (10):

- ?(10) She is remembering that there are infinitely many primes.

On the salient reading of “remember”, (10) is deviant, but it might correctly be used to say that she is in the process of recalling that there are infinitely many primes (see Vendler 1967, p. 104 for more on the linguistic marks of statives).

Third, an FMSO ascribes an attitude to a proposition to the subject. Thus “S  $\Phi$ s that A” entails “S grasps the proposition that A”. To know that there are infinitely many primes, one must grasp the proposition that there are infinitely many primes, so “know” passes the test. A verb with a sense like “is responsible for its being the case that” would fail it. Thus, given that “see” and “remember” are FMSOs, one can see that Olga is playing chess or remember that she was playing chess only if one has a concept of chess. This is not to deny that one’s perceptions and memories may have a content which one lacks the concepts to express; the point is just that the English constructions “see that A” and “remember that A” do not ascribe such content. Other constructions with those verbs behave differently; one does not need a concept of chess to see or remember Olga playing chess.

Fourth, an FMSO is semantically unanalysable. An artificial verb stipulated to mean the same as “believe truly” would not be an FMSO. A semantically analysable expression has a more complex semantic role than that of simply denoting an attitude; its proper treatment would require an account of the meanings from which its meaning is composed. Thus it is best at this stage to concentrate on semantically unanalysable expressions. Verbs such as “know” and “remember” will be assumed to be semantically unanalysable. However, an FMSO is not required to be syntactically unanalysable. In English and some other languages, for example, the addition of the auxiliary “can” often forms an FMSO (Vendler 1967, pp. 104–6). Consider the following pair:

- (11) She felt that the bone was broken.
- (12) She could feel that the bone was broken.

The “could” in (12) is not the “could” of ability; (12) does not mean anything like:

- (13) She had the ability to feel that the bone was broken.

A rough paraphrase of the salient reading of (11) would be: "She intuitively believed that the bone was broken." A rough paraphrase of the salient reading of (12) would be: "She knew by the sense of touch that the bone was broken". Sentence (12), unlike (11), entails "The bone was broken." Thus "could feel" differs from "felt" in two ways: it is factive, and it is perceptual. Neither of these differences would occur if "could feel" were semantically analysable into "could" and "feel", for that would assimilate "could feel" to "had the ability to feel", which is neither factive nor perceptual. "Could feel" is semantically fused. It is an FMSO; "feel" is not.

"Hear" is like "feel" in this respect. Consider:

- (14) She heard that the volcano was erupting.  
 (15) She could hear that the volcano was erupting.

A rough paraphrase of the salient reading of (14) would be: "She heard a report that the volcano was erupting." A rough paraphrase of the salient reading of (15) would be: "She knew by the sense of hearing that the volcano was erupting." Sentence (15), unlike (14), entails "The volcano was erupting." Thus "could hear" differs from "heard" in two ways: it is factive, and it is more directly perceptual. Neither of these differences would occur if "could hear" were semantically a compound of "could" and "hear". "Could hear" is an FMSO; "hear" is not.

"Could see" differs from "see" in only one of the two ways. Consider:

- (16) She saw that the stock market had crashed.  
 (17) She could see that the stock market had crashed.

Both (16) and (17) entail "The stock market had crashed"; there is no difference in factiveness. However, they are naturally read in such a way that (16) would be true and (17) false if she simply saw a newspaper report of the crash; (17) might be true if she saw investors lining the window ledges. In such cases, one could insert "the news" before "that" in (16) but not in (17) – not even

when she has inferred the crash from newspaper reports of other events. In this way, "could see" is more directly perceptual than "saw". This does not prevent both from being FMSOs.

The notion of an FMSO should by now be clear enough to be workable; it can be projected onto new cases. Moreover, it has been explained without essential reference to the notion of knowing, although "know" is an example of an FMSO. It will now be proposed that "know" has a special place in the class of FMSOs.

The proposal is that if  $\Phi$  is any FMSO, then "S  $\Phi$ s that A" entails "S knows that A". If you see that it is raining, then you know that it is raining. If you remember that it was raining, then you know that it was raining. Such entailments are plausible but not uncontroversial (see Unger 1972 and 1975, pp. 158–83 for useful discussion).

It is sometimes alleged that one can perceive or remember that A without knowing that A, because one fails to believe or to be justified in believing that A. Other evidence may give one reason to think that one is only hallucinating what one is in fact perceiving, or only imagining what one is in fact remembering. One abandons the belief, or retains it without justification; either way, it is alleged, one fails to know (Steup 1992 is a recent example of such a view). However, such cases put more pressure on the link between knowing and believing or having justification than they do on the link between perceiving or remembering and knowing. If you really do see *that* it is raining, which is not simply to see the rain, then you know that it is raining; seeing that A is a way of knowing that A. You may not know that you see that it is raining, and consequently may not know that you know that it is raining, but neither condition is necessary for knowing that it is raining. Similarly, if you really do remember *that* it was raining, which is not simply to remember the rain, then you know that it was raining; remembering that A is a way of knowing that A. You may not know that you remember that it was raining, and consequently may not know that you know that it was raining, but neither condition is necessary for knowing that it is raining. But it is far from obvious that you do see or remember that it is or was raining in the cases at issue, and an account will now be suggested on which you do not.

There is a distinction between seeing that A and seeing a situation in which A. One difference is that only the former requires the perceiver to grasp the proposition that A. A normal observer in normal conditions who has no concept of chess can see a situation in which Olga is playing chess, by looking in the right direction, but cannot see *that* Olga is playing chess, because he does not know what he sees to be a situation in which Olga is playing chess. The present cases suggest another difference between the two notions of seeing. By looking in the right direction, you can see a situation in which it is raining. In the imagined case, moreover, you have enough concepts to grasp the proposition that it is raining. Nevertheless, you cannot see *that* it is raining, precisely because you do not know what you see to be a situation in which it is raining (given the unfavourable evidence). On this account, the case is a counterexample to neither the claim that seeing implies knowing nor the claim that knowing implies believing.

Similarly, there is a distinction between remembering that A and remembering a situation in which A. One difference is that only the former requires the rememberer to grasp the proposition that A. Someone whose memory is functioning normally but who has no concept of chess can remember a situation in which Olga was playing chess, but cannot remember *that* Olga was playing chess, because he does not know what he remembers to be a situation in which Olga was playing chess. The present cases suggest another difference between the two notions of remembering. You can remember a situation in which it was raining. In the imagined case, moreover, you have enough concepts to grasp the proposition that it was raining. Nevertheless, you cannot remember *that* it was raining, precisely because you do not know what you remember to be a situation in which it was raining (given the unfavourable evidence). On this account, the case is a counterexample to neither the claim that remembering implies knowing nor the claim that knowing implies believing.

The discussion of FMSOs may be summarized in three principles:

- (18) If  $\Phi$  is an FMSO, from “S  $\Phi$ s that A” one may infer “A”.  
 (19) “Know” is an FMSO.

- (20) If  $\Phi$  is an FMSO, from “S  $\Phi$ s that A” one may infer “S knows that A”.

The latter two principles characterize the concept of knowing uniquely, up to logical equivalence, in terms of the concept of an FMSO. For let “*schnow*” be any term governed by (19') and (20'), the results of substituting “*schnow*” for “know” in (19) and (20) respectively. By (19) and (20'), from “S knows that A” one may infer “S *schnows* that A”. Similarly, by (19') and (20), from “S *schnows* that A” one may infer “S knows that A”. Thus “*schnow*” is logically equivalent to “know”. Note that this argument would fail if (20) held only for *most* FMSOs. In simple terms, “know” is the most general FMSO, the one that applies if any FMSO at all applies.

In the material mode, the claim is that knowing is the most general stative propositional attitude such that, for all propositions  $p$ , necessarily if one has it to  $p$  then  $p$  is true. This is not quite to claim that, for all propositions  $p$ , knowing  $p$  is the most general mental state such that necessarily if one is in it then  $p$  is true. The latter claim fails for necessarily true propositions: every mental state is such that necessarily if one is in it then  $5 + 7 = 12$ , but it does not follow that every mental state is sufficient for *knowing* that  $5 + 7 = 12$ .

It is vital to this account of “know” that “believe truly” does not count as an FMSO. If it did, (20) would permit the invalid inference from “S believes truly that A” to “S knows that A”. The mental state is believing that A, not believing truly that A. To entail knowing, the mental state itself must be sufficient for truth. The condition of semantic unanalysability ensures that “believe truly” does not count as an FMSO.

On this account, the importance of knowing to us becomes as intelligible as the importance of truth. Factive mental states are important to us as states whose essence includes a matching between mind and world, and knowing is important to us as the most general factive stative attitude. Of course, something needs to be said about the nature and significance of this matching, but that is a further problem. Someone who denied that the concept characterized by (18)–(20) is our concept *knows* might even think that it was more useful than the latter.

The states in question are general: different people can be in them at different times. No claim

is made about the essences of their tokens; indeed, the idea of a token state is of doubtful coherence (Steward 1997, pp. 105–34). With respect to general states, the claims of necessity are *de re*, not just *de dicto*. Given that “knowing *p*” rigidly designates a mental state, the *de dicto* claim that the truth of *p* is necessary for knowing *p* implies the *de re* claim that for some mental state *S* the truth of *p* is necessary for *S*.

The account is explicitly not a decomposition of the concept *knows*; if “know” were semantically analysable, it would not be an FMSO. It would certainly be quite implausible to claim that everyone who thinks that John knows that it is raining thereby thinks that John has the most general stative propositional attitude such that, for all propositions *p*, necessarily if one has it to *p* then *p* is true, to the proposition that it is raining. What, then, is the status of the account?

Consider an analogy. Identity is uniquely characterized, up to logical equivalence, by the principles of reflexivity and Leibniz’s Law, just as knowing is uniquely characterized, up to logical equivalence, by (19) and (20). However, it would be quite implausible to claim that everyone who thinks that Istanbul is Constantinople thereby thinks that Istanbul bears to Constantinople the reflexive relation that obeys Leibniz’s Law. The metalogical concepts used in formulating Leibniz’s Law are far more sophisticated than the concepts we use in thinking that Istanbul is Constantinople. In order to have the concept *is* (of identity), one must somehow be disposed to reason according to Leibniz’s Law, but that does not require one to have the metalogical concepts used in formulating Leibniz’s Law. If it did, there would be an obvious danger of an infinite regress. Similarly, in order to have the concept *knows*, one must somehow be disposed to reason according to (18)–(20), but that does not require one to have the metalinguistic concepts used in formulating (18)–(20).

It is no straightforward matter to say what it is for a subject to be disposed to reason according to rules which the subject cannot formulate. Such a subject may even consciously reject the rules; philosophers who mistakenly deny Leibniz’s Law do not thereby cease to understand the “is” of identity. Nevertheless, some such notion does seem to be needed, independently of the account of knowing; the latter account can avail itself of that

notion, whatever exactly it proves to be. The present account of knowing is consistent with the main features of a theory of concepts such as that of Peacocke 1992, on which an account of a concept gives necessary and sufficient conditions for possession of the concept without any need to decompose the concept itself. However, the account is not committed to any general programme of Peacocke’s kind in the theory of concepts.

The present account of knowing makes no use of such concepts as *justified*, *caused*, and *reliable*. Yet knowing seems to be highly sensitive to such factors over wide ranges of cases. Any adequate account of knowing should enable one to understand these connections. This challenge is not limited to the present account: standard accounts of knowing in terms of justification must enable one to understand its sensitivity to causal factors, and standard accounts of knowing in terms of causal factors must enable one to understand its sensitivity to justification; none of these tasks is trivial.

One way for the present account to meet the challenge is by exploiting the metaphysics of states. For example, a form of the essentiality of origins may apply to states; a necessary condition of being in some states may be having entered them in specific ways. States of perceiving and remembering have this feature, requiring entry along a specific kind of causal path. Thus the importance of causal factors in many cases of knowing is quite consistent with this account. More obviously, having an inferential justification of a specific kind may be essential to being in some mental states; having a proof is clearly a factive mental state. Thus the importance of justification in many cases of knowing is equally consistent with this account. Of course, these remarks merely adumbrate a strategy, without carrying it out. ... We can see epistemology as a branch of the philosophy of mind. If we try to leave epistemology out of the philosophy of mind, we arrive at a radically impoverished conception of the nature of mind.

## 5 Knowing and Believing

The account of knowing above makes no essential mention of believing. Formally, it is consistent with many different accounts of the relation

between the two concepts. Historically, however, the view of knowing as a mental state has been associated with the view that knowing entails *not* believing. Prichard is a case in point (1950, pp. 86–8). On standard analyses of knowing, in contrast, knowing entails believing. On some intermediate views, knowing is consistent both with believing and with not believing. It is therefore natural to ask how far the present account of knowing constrains the relation between knowing and believing.

We have two schemas to consider:

- (21) If S knows that A then S believes that A.  
 (22) If S knows that A then S does not believe that A.

If (21) is invalid, then the programme of analysing the concept *knows* as a conjunction of *believes* with *true* and other concepts is stillborn. Once the programme has been abandoned, (21) can be examined without prior need for its vindication.

The schema (22) is quite implausible. Whether I know that A on being told that A depends constitutively on whether my informant knew that A (amongst other factors). Whether I believe that A on being told that A does not depend constitutively on whether my informant knew that A; it would have to if knowing excluded believing. Of course, when one can describe someone as knowing that A, it is conversationally misleading simply to describe her as believing that A, but that is not to say that it is false. Not all believing is mere believing. We should reject (22).

The schema (21) does not sound *trivially* valid, as the schema “If S knows that A then A” does. When the unconfident examinee, taking herself to be guessing, reliably gives correct dates as a result of forgotten history lessons, it is not an obvious misuse of English to classify her as knowing that the battle of Agincourt was in 1415 without believing that it was. But intuitions differ over such cases; it is not very clear whether she knows and not very clear whether she believes. In a case in which she was taught incorrect dates and repeats them with equal unconfidence, she is in an at least somewhat belief-like state, which she is also in when she was taught the correct dates. We have no clear counterexamples to (21) (see Radford 1966, Armstrong 1973, pp. 138–49, and Shope 1983, pp. 178–87 for further discussion of such cases).

There is a wide grammatical divergence between the verbs “know” and “believe” not suggestive of closely connected terms. For example, in a context in which I have predicted that it will rain, “You know what I predicted” has a reading on which it is true if and only if you know that I predicted that it will rain, whereas “You believe what I predicted” has no reading on which it is true if and only if you believe that I predicted that it will rain. There are many further grammatical differences between “know” and “believe” (see Austin 1946, Vendler 1972, pp. 89–119, and Shope 1983, pp. 171–8, 191–2). One explanation of such facts, proposed by Vendler, is that “know” and “believe” take different objects: what one knows is a fact, what one believes a proposition, where a fact is not a true proposition. A contingently true proposition, unlike a contingent fact, could have been false and still have existed. If so, then knowing is not a *propositional* attitude, and much of the terminology of this book might need revision, although the substance of the account would remain. Vendler’s explanation makes it hard to see why (21) should be valid. However, it is not strictly inconsistent with the validity of (21), since “that A” may refer to a fact in the antecedent and to a proposition in the consequent.

If “that A” refers to a fact in the context “S knows that A”, then we might expect “that A” to suffer reference failure when “A” is false. Consequently, we might expect “S knows that A” and “S does not know that A” not to express propositions. But if “A” is false, “S knows that A” expresses a false proposition and “S does not know that A” a true one. Perhaps we could treat “that A” as elliptical for “the fact that A” and analyse it by a Russellian theory of definite descriptions. The reference of “fact that A” in the definite description is presumably determined by the proposition  $p$  expressed by “A”; it is therefore some function  $f$  of  $p$ . Thus to know that A is to know the  $f(p)$ , and hence to stand in a complex relation expressed by “know”, “the”, and “ $f$ ” to the proposition expressed by “A”. But then with only a slight change of meaning we could use the word “know” for that complex relation to a proposition. Thus, even on a view like Vendler’s, knowing would still involve a propositional attitude. However, it is very doubtful that there are any such things as facts other than true propositions

(see Williamson 1999 for an argument). Moreover, the propriety of remarks like “I always believed that you were a good friend; now I know it” and “Long before I knew those things about you I believed them” suggest that “believe” and “know” do take the same kind of object. Vendler’s account is not accepted here.

The present account of knowing might be thought inconsistent with the validity of (21), on the grounds that it provides no basis for a conceptual connection between believing and knowing. That would be too quick. Section 3 already noted that not every conceptually necessary condition is a conjunct of a conjunctive analysis. It is a mistake to assume that (21) is valid only if that connection is explicable by an analysis of *knows* in terms of *believes*. Consider an analogy: it may be a priori that being crimson is sufficient for being red, but that implication need not be explained by an analysis of one colour concept in terms of the other. One can grasp either concept without grasping the other, by being shown examples of its application and non-application. Neither concept relies on the other in demarcating conceptual space. Nevertheless, the area demarcated by one concept might be so safely within the area demarcated by the other that one could know by a priori reflection that the former is sufficient for the latter. Similarly, the area demarcated by the concept *knows* might be so safely within the area demarcated by the concept *believes* that one could know (21) by a priori reflection. That is quite consistent with, although not entailed by, the account of knowing in section 4.

An alternative proposal is to reverse the direction of analysis, and validate (21) by an analysis of *believes* in terms of *knows*. The simplest suggestion is that the concept *believes* is analysable as a disjunction of *knows* with other concepts. The word “opine” will be used here as a term of art for the rest of the disjunction. On this analysis, one believes *p* if and only if one either knows *p* or opines *p*. Given that opining *p* is incompatible with knowing *p*, it follows that one opines *p* if and only if one believes *p* without knowing *p*. A similar view has been proposed by John McDowell (1982), building on the disjunctive account of perceptual experience developed by J. M. Hinton (1967 and 1973) and Paul Snowdon (1980–1 and 1990; see also Child 1994, pp. 143–64,

Dancy 1995, and Martin 1997). In McDowell’s terminology, believing is not the highest common factor of knowing and opining. There is no such common factor. Rather, knowing and opining are radically different, mutually exclusive states, although instances of the latter are easily mistaken for instances of the former. Given a distinction between facts and true propositions, one could contrast knowing and opining somewhat as Vendler contrasts knowing and believing: to know is to be acquainted with a fact; to opine is to be acquainted with no more than a proposition. But the disjunctive conception does not require such an ontology of facts.

Not all those who advocate a disjunctive conception would claim that it provides a conceptual analysis. That claim faces difficulties additional to the generally dim prospects for conceptual analysis evoked in section 3. If the concept *believes* is the disjunction of *knows* and *opines*, then it must be possible to grasp the concept *opines* without previously grasping the concept *believes*. For otherwise, since grasping a disjunction involves grasping its disjuncts, it would be impossible to grasp the concept *opines* for the first time. Now “opine” was introduced as a term of art; how is it to be explained? The natural explanation is that to opine a proposition *p* is to have a mere belief *p*, which is presumably to believe *p* without knowing *p*, but that explanation uses the concept *believes*. It does not permit one to grasp *opines* without already grasping *believes*. The explanation that to opine *p* is to be of the opinion *p* does no better, for “be of the opinion” as ordinarily understood is just a rough synonym of “believe”. In particular, once it is conceded – as it is by the disjunctive conception – that “know” implies “believe”, little reason remains to deny that “know” implies “be of the opinion”, too.

Can we explain “opine” in terms of “know”? A first attempt is this: one opines the proposition *p* if and only if one is in a state which one cannot discriminate from knowing *p*, in other words, a state which is, for all one knows, knowing *p*. That cannot be quite right, for if one cannot grasp the proposition *p* then one cannot discriminate one’s state from knowing *p*; but one does not believe *p*, and therefore does not opine it. To avoid that problem, we can revise the definition thus: one opines *p* if and only if one has an attitude to the proposition *p* which one cannot discriminate

from knowing, in other words, an attitude to  $p$  which is, for all one knows, knowing. However, that definition does not help a *disjunctive* analysis of believing. For if one knows  $p$ , then trivially one has an attitude to  $p$  which one cannot discriminate from knowing; one cannot discriminate something from itself. Thus the first disjunct, “One knows  $p$ ”, entails the second disjunct, “One opines  $p$ ”. The whole disjunction would therefore be equivalent to its second disjunct, and the disjunctive form of the definiens would be a mere artefact of conceptual redundancy. To tack the qualification “but does not know  $p$ ” onto the end of the definition of “opine” would make no significant difference, for since “One either knows  $p$  or has an attitude to  $p$  which one cannot discriminate from knowing but does not know  $p$ ” is still equivalent to “One has an attitude to  $p$  which one cannot discriminate from knowing  $p$ ”, the disjunctive form would remain a mere artefact.

Alternatively, “opine” might be explained as the disjunction of several more specific disjuncts, such as “be under the illusion”, “be irrationally certain” and so on. However, it is very doubtful that, without using the concept *believes*, one could extend such a list to include all the different ways in which someone can believe without knowing. Those ways seem to be indefinitely various. How could one even specify, without using the concept *believes*, all the states in which someone can believe  $p$  falsely? If the list of disjuncts is open-ended, one could not grasp how to go on without realizing that one must list the ways in which someone can believe without knowing. Thus the explanation of “opine” illicitly relies on a prior grasp of the concept *believes*.

The phenomenon just noted also threatens more metaphysical disjunctive accounts which do not attempt conceptual analysis, instead making their claims only about the underlying facts in virtue of which the concepts apply. Such an account of believing might deny that believing is itself a unified state, insisting that it is necessary but not a priori that one believes  $p$  if and only if one is in either the state of knowing  $p$  or the state of opining  $p$ . Since conceptual analysis is no longer in question, the replacement of “opining” by “merely believing” is not objectionable on grounds of circularity. The trouble is rather that there is no *more* reason to regard merely believing  $p$  as a unified mental state than to regard believ-

ing  $p$  as such. What unifies Gettier cases with cases of unjustified false belief is simply that in both, the subject believes without knowing; a good taxonomy of believing would not classify them together on the basis of some positive feature that excludes knowing. Moreover, it is hard to see how such a taxonomy could describe every species of believing without using the concept *believes*. But if a good taxonomy of believing does use the concept *believes*, that undermines the denial that believing is a unified state. Similar objections apply to disjunctive accounts of perception, appearance, and experience. For example, there is no reason to postulate a unified mental state equivalent to its appearing to one that  $A$  while one does not perceive that  $A$ .

A strictly disjunctive account of belief is not correct at either the conceptual or the metaphysical level. However, the disjunctive account was brought into play as a simple means to reconcile the account of knowing in section 4 with the supposed validity of (21) (knowing entails believing). There are other means to that end. A non-disjunctive analysis of *believes* might also validate (21). For example, (21) is a corollary of an analysis of *believes* itself on the lines of the definition of *opines* above: one believes  $p$  if and only if one has an attitude to the proposition  $p$  which one cannot discriminate from knowing, in other words, an attitude to  $p$  which is, for all one knows, knowing. That definition suggestively makes knowing central to the account of believing. One attraction of such an account is that it opens the prospect of explaining the difficulty, remarked by Hume, of believing  $p$  at will in terms of the difficulty of knowing  $p$  at will. The analysis is also consistent with the account of knowing in section 4.

Although that analysis provides a reasonable approximation to our concept *believes*, it does not fully capture the concept. It incorrectly classifies as believing that food is present a primitive creature which lacks any concept of knowing and merely desires that food is present; for all the creature knows, its attitude to the proposition that food is present is knowing. Equally incorrectly, the account classifies as not believing that there is a god someone who consciously takes a leap of faith, knowing that she does not know that there is a god. Both examples, however, are compatible with the variant idea that to believe  $p$  is to treat  $p$  as if one knew  $p$  – that is, to treat  $p$  in

ways similar to the ways in which subjects treat propositions which they know. In particular, a factive propositional attitude to a proposition is characteristically associated with reliance on it as a premise in practical reasoning, for good functional reasons; such reliance is crucial to belief. A creature which lacks a concept of knowing can still treat a proposition in ways in which it treats propositions which it knows. The primitive creature does not treat the proposition that food is present like that when merely desiring that food is present; it does not use the proposition as a premise in practical reasoning. By contrast, the person who genuinely believes that there is a god by a leap of faith does rely on that premise in such reasoning. The unconfident examinee who tentatively gives *p* as an answer is little disposed to rely on *p* as a premise, and for that reason does not clearly believe *p*, but for the same reason does not clearly know *p*. Although a full-blown exact conceptual analysis of *believes* in terms of *knows* is too much to expect, we can still postulate a looser connection along these lines.

If believing *p* is, roughly, treating *p* as if one knew *p*, then knowing is in that sense central to believing. Knowledge sets the standard of appropriateness for belief. That does not imply that all cases of knowing are paradigmatic cases of believing, for one might know *p* while in a sense treating *p* as if one did not know *p* – that is, while treating *p* in ways untypical of those in which subjects treat what they know. Nevertheless, as a crude generalization, the further one is from knowing *p*, the less appropriate it is to believe *p*. Knowing is in that sense the best kind of believing. Mere believing is a kind of botched knowing.<sup>8</sup> In short, belief aims at knowledge (not just truth). ...

Although the letter of disjunctive accounts has been rejected, the spirit may have been retained. For on the account in section 4, believing is not the highest common factor of knowing and mere believing, simply because it is not a factor of knowing at all (whether or not it is a necessary condition). Since that point is consistent with the claim that believing is common to knowing and mere believing, the claim is harmless. It no more makes the difference between knowing and mere believing extrinsic to a state than the point that continuity is common to straight and curved lines makes the difference between straight and curved extrinsic to a line. To know is not merely to believe while various other conditions are met; it is to be in a new kind of state, a factive one. What matters is not acceptance of a disjunctive account of believing but rejection of a conjunctive account of knowing.<sup>9</sup> Furthermore, the claim that belief is what aims at knowledge is consonant with the suggestion in disjunctive accounts that illusion is somehow parasitic on veridical perception. Properly developed, the insight behind disjunctive theories leads to a non-conjunctive account of knowledge and a non-disjunctive account of belief.

While belief aims at knowledge, various mental processes aim at more specific factive mental states. Perception aims at perceiving that something is so; memory aims at remembering that something is so. Since knowing is the most general factive state, all such processes aim at kinds of knowledge. If a creature could not engage in such processes without some capacity for success, we may conjecture that nothing could have a mind without having a capacity for knowledge.

## Notes

- 1 McDowell 1995 and Gibbons 1998 defend closely related conceptions of knowing as a mental state. See also Guttenplan 1994 and Peacocke 1999, pp. 52–5.
- 2 See Shope 1983 for the history of a decade of research into the analysis of knowing after Gettier 1963; an equally complex book could be written on post-1983 developments. Not all this work aims to provide an analysis in the traditional sense; see Shope 1983, pp. 34–44.
- 3 Craig 1990 makes an interesting attempt to explain the point of the concept of knowledge in the light of the failure of analyses of the standard kind. However, on the present view it remains too close to the traditional programme, for it takes as its starting point our need for true beliefs about our environment (1990: 11), as though this were somehow more basic than our need for knowledge of our environment. It is no reply that believing

truly is as useful as knowing, for it is agreed that the starting point should be more specific than “useful mental state”; why should it be specific in the manner of “believing truly” rather than in that of “knowing”? ...

- 4 For sophisticated but unconvincing defence of conceptual analysis see Jackson 1998 and Smith 1994, pp. 29–56, 161–4. However, the kind of analysis they defend constitutes little threat to the claim that knowing is a mental state in every reasonable sense of the latter term. They provide no reason to suppose that the concept *knows* can be non-trivially analysed in any sense in which paradigmatic mental concepts cannot be, or that it is somehow posterior in the order of analysis to the concept *believes*. See also Fodor 1998 for a discussion of the demise of definition.
- 5 We must also assume Russell’s conception of propositions as at the level of reference rather than sense. In effect, Evans 1982 combines the Principle of Acquaintance with a conception of acquaintance much less extreme than Russell’s. Of course, Russell’s extremism here is no mere extraneous dogma; it is an attempt to solve puzzles about the identity and non-existence of denotation in intentional contexts.

Unfortunately, the cure is worse than the disease.

- 6 As noted in the Introduction, we cannot define “*x* is a parent of *y*” by “*x* is an ancestor of *y* and *x* is not an ancestor of an ancestor of *y*”.
- 7 A further ground for suspicion of analyses of the concept *knows* in terms of the concept *believes* is that they seem to imply that the latter concept is acquired before the former. Data on child development suggest, if anything, the reverse order (see Perner 1993, pp. 145–203 for discussion of relevant work). Crudely: children understand ignorance before they understand error. Naturally, the data can be interpreted in various ways, and their bearing on the order of analysis depends on subtle issues in the theory of concepts.
- 8 See also Peacocke 1999, p. 34.
- 9 Martin 1997, pp. 88–90 questions whether a parallel account of perception and appearance will serve the purposes of naive realism, on the grounds that it does not entail the naive realist’s distinctive claims about the phenomenology of perception. But a parallel account in terms of a factive mental state of conscious perceptual awareness may capture such claims.

## References

- Armstrong, D. M. 1973. *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press).
- Austin, J. L. 1946. “Other minds.” *Proceedings of the Aristotelian Society*, supp. 20, pp. 148–87.
- Child, T. W. 1994. *Causality, Interpretation and the Mind* (Oxford: Clarendon Press).
- Craig, E. J. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press).
- Dancy, J. 1995. “Arguments from illusion.” *Philosophical Quarterly* 45, pp. 421–38.
- Evans, M. G. J. 1982. *The Varieties of Reference*, ed. J. H. McDowell (Oxford: Clarendon Press).
- Fodor, J. A. 1998. *Concepts: Where Cognitive Science Went Wrong* (Oxford: Clarendon Press).
- Gettier, E. 1963. “Is justified true belief knowledge?” *Analysis*, 23, pp. 121–3.
- Gibbons, J. 1998. “Truth in action.” Unpublished typescript.
- Grice, H. P. 1989. *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).
- Guttenplan, S. 1994. “Belief, knowledge and the origins of content.” *Dialectica* 48, pp. 287–305.
- Hintikka, K. J. J. 1962. *Knowledge and Belief* (Ithaca, NY: Cornell University Press).
- Hinton, J. M. 1967. “Visual experiences.” *Mind* 76, pp. 217–27.
- Howson, C. 1996. “Epistemic probability.” *Proceedings of the Aristotelian Society*, supp. 70, pp. 63–77.
- Jackson, F. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (Oxford: Oxford University Press).
- McDowell, J. H. 1982. “Criteria, defeasibility, and knowledge.” *Proceedings of the British Academy* 68, pp. 455–79. Partly reprinted with revisions in Dancy (ed.) 1988. *Perceptual Knowledge* (Oxford: Oxford University Press).

- 1995. "Knowledge and the internal." *Philosophy and Phenomenological Research* 55, pp. 877–93.
- Martin, M. G. F. 1997. "The reality of appearances." In R. M. Sainsbury (ed.) *Thought and Ontology* (Milan: FrancoAngeli).
- Peacocke, C. A. B. 1992. *A Study of Concepts* (Cambridge, MA: MIT Press).
- 1999. *Being Known* (Oxford: Clarendon Press).
- Perner, J. 1993. "Fitch and intuitionistic knowability." *Analysis* 50, pp. 182–7.
- Prichard, H. A. 1950. *Knowledge and Perception* (Oxford: Clarendon Press).
- Radford, C. 1966. "Knowledge – by examples." *Analysis* 27, pp. 1–11.
- Russell, J. B. A. W. 1910–11. "Knowledge by acquaintance and knowledge by description." *Proceedings of the Aristotelian Society* 11, pp. 108–28. Page reference to reprinting in Salmon and Soames 1988.
- Salmon, N. U. and Soames, S. (eds) 1988. *Propositions and Attitudes* (Oxford: Oxford University Press).
- Shope, R. K. 1983. *The Analysis of Knowing: A Decade of Research* (Princeton: Princeton University Press).
- Smith, M. 1994. *The Moral Problem* (Oxford: Blackwell).
- Snowdon, P. 1980–1. "Perception, vision and causation." *Proceedings of the Aristotelian Society* 81, pp. 175–92.
- 1990. "The objects of perceptual experience." *Proceedings of Aristotelian Society*, supp. 64, pp. 121–50.
- Steward, H. 1997. *The Ontology of Mind: Events, Processes and States* (Oxford: Clarendon Press).
- Unger, P. 1972. "Propositional verbs and knowledge." *Journal of Philosophy* 69, pp. 301–12.
- 1975. *Ignorance: A Case for Scepticism* (Oxford: Oxford University Press).
- Vendler, Z. 1967. *Linguistics in Philosophy* (Ithaca, NY: Cornell University Press).
- 1972. *Res Cogitans* (Ithaca, NY: Cornell University Press).
- Williamson, T. 1999. "Truthmakers and the Converse Barcan Formula." *Dialectica* 53, pp. 253–70.