

CHAPTER THREE



The Concept of Knowledge

Having examined Descartes's epistemological view as a kind of prologue, we will now turn to a more detailed consideration of a variety of more specific epistemological issues, focusing mainly on those that naturally arise out of his discussion. Our first specific concern will be to achieve a deeper understanding of the concept of *knowledge* itself. What is it to *know* something? What, that is, are we saying of a person when we ascribe knowledge to him or her? A further set of questions, already briefly noticed in chapter 1, concerns the significance of the concept of knowledge. *Why* does it involve the specific conditions that it does? How do those conditions fit together or connect with each other in an intelligible way? And, most fundamentally, why do or should we *care* about knowledge? Why is having knowledge important and valuable in the way that we normally take it to be (if indeed it really is)?

We have already encountered one specific account of the concept of knowledge, the one that Descartes seems to have roughly in mind (though without formulating it very explicitly) in the *Meditations*. According to that account, for a person S to know some proposition¹ P at some time t, the following three conditions must be satisfied (with the subscripts indicating that these are the conditions of the *Cartesian* conception of knowledge):

- 1_C. S must *believe* or *accept* P at t without any doubt.
- 2_C. P must be *true*.
- 3_C. S must have at t a *reason* or *justification* that *guarantees* that P is true.²

It is obvious on reflection that condition (3_C) makes condition (2_C) redundant and so unnecessary: If S has a reason that *guarantees* P's truth, then it follows automatically that P is in fact true. But since there are other accounts of knowledge that we will want to consider and compare with this one in which the condition that is parallel to Descartes's condition (3_C) does not in this way entail that the condition parallel (and in fact usually identical) to (2_C) is satisfied, it will be clearer to list condition (2_C) separately in spite of its redundancy in this case.

The Cartesian account of knowledge is in fact one specific version of a more general account of knowledge that has come to be generally referred to as "the traditional conception of knowledge." According to this more general account, knowledge requires the satisfaction of three conditions at least roughly parallel to Descartes's: (1) a belief or acceptance condition, (2) a truth condition, and (3) a reason or justification condition (so that accounts of this kind are often referred to as *justified true belief* accounts or definitions of knowledge). Other specific versions of this general account almost always share Descartes's truth condition (2_C), but differ somewhat in their specification of the belief or acceptance condition parallel to (1_C) and to a wider and more serious extent in their specification of the reason or justification condition parallel to (3_C). As we will eventually see, there are also many recent accounts that add a further condition (4), while still retaining conditions parallel to Descartes's three.

We will begin by examining the three general kinds of conditions included in the traditional account, considering the rationale for including each of them as an essential part of the concept of knowledge, trying to understand the general nature of each condition, and discussing, briefly for condition (1) and more extensively for condition (3), some of the different ways in which the general condition in question has been further specified by different versions of the traditional account.

The Belief or Acceptance Condition

The basic rationale for this first general condition is quite straightforward: Someone who is in serious doubt as to whether a particular proposition is true or, perhaps even more obviously, who has never so much as considered or entertained that proposition can surely not be correctly said to have *knowledge* of it. If I am completely uncertain about whether it will rain tomorrow, then I do not know that it will (or that it won't). And if it has never so much as occurred to me that my roof might be leaking, then again I plainly do not know that it is—even if in fact it is leaking and even if I have what would

be good evidence for this being so if I were to recognize it as such (there are damp spots on the rug and distinctive streaks on the walls).

The most obvious way to satisfy a condition of this general sort would be for the person in question to be in the conscious state of explicitly considering and assenting to the proposition in question. This might involve, as the formulation just given seems to suggest, a two-stage process: for example, my wife suggests to me that perhaps the roof is leaking, and after considering the evidence, I end up becoming convinced that this is indeed what is going on. It is also possible, however, that the truth of the proposition strikes me as obvious as soon as it enters my mind, without any preliminary stage of consideration.

But while this is one way in which a condition of the indicated sort might be satisfied, it seems reasonably clear that it is not the only way, that people can and do know many things at a particular time that they do not have explicitly in mind at that time. I am about to give you an example of such a piece of knowledge, something that it seems plain that you in fact know even as you read these words though you do not at the moment have it explicitly in mind. (It was, of course, precisely to produce this situation that I didn't initially specify the proposition in question.) Consider, then, the claim that you are a human being, where what is intended is that each reader formulate the appropriate version of this general sort of claim, the one that applies to himself or herself. My suggestion is that the claim in question is something that you knew to be true while reading the earlier part of the present paragraph, even though you almost certainly did not have it explicitly in mind. If this is right, then a correct formulation of the belief or acceptance condition for knowledge should not require explicit, conscious acceptance of the relevant proposition at the time in question, even though this is clearly *one* way in which such a condition might be satisfied.

Perhaps the most standard way of handling this point is to first formulate the condition in question as the requirement that the person who has knowledge *believe* the proposition in question and to then distinguish two different kinds of belief (or, as it is often put, two "senses" of the term "belief"): *occurrent* belief, which is what happens when the person has the proposition explicitly in mind and accepts or assents to it; and *dispositional* belief, where the person does not have the proposition explicitly in mind, but is *disposed* to accept or assent to it, that is, *would* accept or assent to it if the issue were raised. Thus in the case of the example just given, the suggestion would be that each of the readers of this book had a dispositional belief (or believed dispositionally) that he or she is a human being, though it is quite possible that none of you had an occurrent belief to this effect at the time just before the proposition was explicitly mentioned.

There is, however, a problem lurking here which needs to be dealt with. The following situation sometimes, perhaps even fairly frequently, occurs: there is a proposition that a person has *never* consciously or explicitly considered, still less consciously assented to, but which is in some way obvious enough that he or she would immediately accept or assent to it if it were proposed. In such a case, the requirements of the so-called dispositional sense of “belief,” as just given, seem to be satisfied, but it still seems plainly wrong to say that the person believes the proposition in question—and even more plainly wrong to say that he or she knows it, even if all of the other conditions for knowledge should happen to be satisfied.³ Consider as an example here the version of the leaking roof case discussed earlier in which I would accept at once the proposition that my roof is leaking if proposed by my wife or if it just happened to occur to me, but where I neither have it presently in mind nor have accepted or assented to it earlier, and suppose also that Descartes’s conditions are at least roughly on the right track. In such a situation, even if the evidence I have (the wet spots on the floor and distinctive streaks on the wall) would be enough to satisfy the correct version of the reason or justification condition, and even if the claim in question is in fact true (and even if whatever further conditions there might be in addition to or instead of these two are also satisfied), it still seems plainly wrong to say that I *know* that my roof is leaking—even though it does not seem wrong to say of the readers of this book that they knew that they were human beings even at the time prior to my explicitly suggesting this proposition to them.

For this reason, rather than defining dispositional belief in the way suggested earlier, it should be specified instead as the dispositional state in which (a) one has *previously* explicitly considered and consciously accepted or assented to the proposition in question, and (b) as a direct result of this prior acceptance or assent, would accept or assent to it again if the question were explicitly raised. It is also perhaps a bit clearer not to use the term “belief” for the first alternative of conscious or explicit acceptance or assent. We can then say that condition (1) of the standard conception of knowledge should be understood to require that the person in question *either* explicitly and consciously accepts *or* else (dispositionally) believes the proposition in question at the time in question. (Having clarified the point, I will sometimes for the sake of brevity follow fairly standard philosophical practice by using the term “belief” to cover both dispositional belief under the corrected specification and conscious, explicit acceptance or assent.)

One further issue, the main one that distinguishes different versions of the belief or acceptance condition, is how *strongly* the person must accept or believe the proposition in question, that is, how strongly they must

be convinced that it is true. The Cartesian view, as formulated earlier, requires that the person have no doubt at all that the proposition is true, a condition that is also sometimes formulated by saying that he must be *certain* of it. This is a very strong version of the belief or acceptance requirement—one that many or probably most of the things that we seem ordinarily to regard as instances of knowledge (see the list of examples in chapter 1) would not satisfy, though this fact is obscured somewhat by a tendency to exaggerate when saying that claims are certain or indubitable. Thus I might well say that I am certain that my dog is in the yard where I just left him or certain that Obama was elected president in 2008 or certain that it is very hot in the center of the sun; but if pressed, I would have to admit that none of these claims is really beyond all possible doubt. (Try to think in each of these cases and in various others of ways in they might be false that you can see to be at least *possible*. Do so, if you can, without appealing to anything as outlandish as the Cartesian evil genius.)

A significantly weaker version of the belief or acceptance condition would say instead merely that the person must be fairly confident, reasonably sure in his or her belief or acceptance of the proposition in question. This is a requirement that seems to agree much better with our common-sense judgments about the extent of our knowledge (as reflected in the list in chapter 1). Is this a good reason for thinking that it is this second, weaker requirement that is in fact correct as one part of an account of the concept of knowledge?

A Digression on Method

There is an important—and difficult—issue of philosophical method pertaining to this last point, one that is indeed also relevant to the earlier examples, and this is as good a place as any to discuss it. As we have already seen, there are many claims of many different kinds, roughly indicated in the earlier list, that from a “common-sense” or “intuitive” standpoint count as cases of knowledge (with obviously some significant variation from person to person). What this means is that most ordinary people and even most philosophers, if asked to consider whether such an example is a case of knowledge, would be inclined to say without much hesitation that it is. Thus we have (a) a proposed requirement for knowledge, the Cartesian requirement that a proposition that is known must be believed without any doubt or with certainty, together with (b) a large number of commonsensically or intuitively accepted cases of knowledge that do not satisfy the requirement in question—and so, if this requirement is correct, must not be genuine. Either

the intuitive judgments about these particular cases or the proposed requirement must apparently be mistaken (assuming that the concept of knowledge is unambiguous), but how are we to decide between these two alternatives?

There are at least two reasonably clear things to be said about this issue, though they are not, alas, sufficient to resolve it. First, common-sense or intuitive judgments about particular cases are a central and essential part of our basis for understanding and delineating concepts like the concept of knowledge. This is just to say that if all such judgments were dismissed as undependable, we would have little handle left on such concepts. (Imagine trying to figure out what knowledge is if you have no idea at all which particular examples in fact qualify as cases of knowledge. How would you begin? What would you rely on?)

But, second, while common-sense or intuitive judgments of the sort in question are in this way indispensable, there is no apparent reason to regard them as somehow simply incapable of being mistaken. This would be so even if it were not the case—as in fact it is—that the common-sense or intuitive judgments of different people or of the same person at different times often conflict with each other. And if it is possible for such judgments to sometimes be mistaken, then it is hard to rule out completely the possibility that they might be largely or even entirely mistaken, so that some requirement that most or all of the intuitive or commonsensical cases of alleged knowledge fail to satisfy might still be correct.

The upshot of these considerations is thus rather inconclusive. It seems right to say that the fact that the Cartesian version of the first condition conflicts with our common-sense or intuitive judgments about cases of knowledge counts against it and in favor of the weaker version mentioned above. (Similarly, on an earlier issue, the fact that it seems intuitively wrong to say that I know that my roof is leaking when the proposition in question has never explicitly occurred to me counts in favor of a version of the first condition that would not be satisfied in that case.) But this resolution of the issue between the two versions of the first condition is not decisive, since there is no guarantee that the relevant “intuitions” are correct. At least some possibility remains that the Cartesian condition is correct after all—in which case, hopefully, we might be able to find further reasons of some sort that point in this direction.

The Truth Condition

The rationale for the truth condition is simply that one cannot know what is not the case, something that almost no philosopher has seriously disputed. If

I know that my car is in the parking lot, then it must actually be there; if it is not, then I did not in fact *know* that it was there, no matter how sure I may have been and how strong my reasons or justification may have been.

One thing that sometimes makes people balk at accepting the truth condition is that someone can, of course, *think* that he knows something when in fact it is not true. Thus in the case just given, I may still think that I know that my car is in the parking lot. Similarly, many people living prior to the exploits of Columbus believed that the earth was flat and thought that this was something that they knew. And many scientists and others living prior the work of Einstein believed that Newtonian mechanics was an exactly correct description of the behavior of material bodies and again thought that this was a case of knowledge, indeed an exceptionally clear one. Moreover, in describing cases of this kind, it is sometimes tempting, and perhaps even useful in some ways, to temporarily take the point of view of the people in question and thus describe the situation by saying that they knew the claim in question—that is, that from their perspective it clearly seemed that they knew. According to all versions of the traditional conception of knowledge, however, such ascriptions of knowledge where the proposition in question is false are always mistaken, however reasonable and obvious they may have seemed to the people in question.

Here we have a somewhat more subtle example of the appeal to intuitive or common-sense judgments. In general, it seems intuitively wrong to ascribe knowledge where the claim in question is not in fact true. This is why a person who claims to know something will normally withdraw that claim when it is demonstrated in some way that the claim in question is mistaken and will concede that he or she did not know after all. But there are also certain cases, such as that of beliefs about the shape of the earth prior to Columbus, where there seems to be *something* right about saying that the people in question knew something that wasn't so. This conflict is resolved by pointing out that the ascription of knowledge in such cases in effect reflects the point of view of the people in question, from which the proposition *seemed* true; thus this ascription can still be said to be mistaken from a more objective standpoint in which the falsity of the claim is acknowledged.

A related problem that you may perhaps have with the truth condition arises from worrying about how you could ever tell that it is satisfied. As we will see further below, a person does in a way have to determine that the proposition is true, according to the traditional conception at least—something that is accomplished by appeal to the reasons or justification for it. But it is tempting to make the mistake of thinking of the truth condition as one whose satisfaction has to be somehow determined by the would-be knower

independently of the satisfaction of the other two conditions, and the problem is then that there is no apparent way to do this. As the point is sometimes put, you cannot just “step outside” of your own subjective perspective and observe independently that the claim that you believe and for which you perhaps have good reasons or justification is also true—there is just no way to occupy such a “God’s-eye” perspective. But a proponent of the traditional conception will reply that what this shows is not that the truth condition is mistaken, but rather that it is a mistake to think of it as a condition that a person must determine independently to be satisfied in order to have knowledge; instead, it is just a condition that must in fact *be* satisfied (something that is in fact true of all of the conditions in question).

A useful way in which this point is sometimes put is to say that the concept of knowledge is a “success” concept, that is, that it describes the successful outcome of a certain kind of endeavor. The aim of the cognitive enterprise is truth: we want our beliefs to correctly describe the world. And, according to the traditional account of knowledge, we attempt to accomplish this by seeking beliefs for which we have good reasons or strong justification. When this endeavor is successful, that is, when the justified beliefs thus arrived at are in fact also true, then we have knowledge; when it fails, when the resulting strongly justified beliefs are not in fact true,⁴ we have only what might be described as “attempted knowledge.” But the distinction between genuine and merely attempted knowledge is not one that we have to, or indeed in the short run could, independently draw. (A crude but still helpful comparison: When shooting an arrow at a target, the aim is to hit the target, and this is something that we attempt to achieve by aiming carefully. But whether or not we succeed depends on whether the arrow does in fact hit the target, and this may be so, in which case we have succeeded, even if we have no independent way to establish that it is so—even if the target is only briefly visible and cannot, for some reason, be examined later.)

A deeper and more difficult question, one that is rather more metaphysical than epistemological in character, concerns the nature of truth itself: What does truth amount to? What does it mean to say that a particular proposition is true? Here there is one answer that is both the most widely accepted and also the one that is seemingly in accord with common sense. But it is an answer that philosophers of very different persuasions have often regarded as problematic or even as not fully intelligible. Thus we need to take a brief look at this controversy, even though a full discussion of it is beyond the scope of this book.

The widely accepted, commonsensical view is what has come to be known as “the correspondence theory of truth.” It says that a proposition

is true if it *corresponds to* or *agrees with* the relevant aspect or part of reality. Thus, for example, for the proposition that my car is in the parking lot to be true, according to the correspondence theory, is for the content of this proposition (that is, what I believe or accept when I believe or accept this proposition) to agree with or match the appropriate aspect or chunk of independent reality—in this case, the physical configuration that involves a certain complicated structure of metal, plastic, rubber, and so forth (my car) being physically juxtaposed in the right way or not with a certain piece of asphalt (the parking lot). The physical configuration that would make the proposition true is something that one could point to or physically mark off (with police tape or by building a box around it) quite independently of the proposition or the various beliefs that involve it; the various ones that would make it false are in general less localized, but could also be pointed at in a way by indicating the two separate elements and their failure to realize the indicated relation.

Many different sorts of problems and objections have been raised in relation to the correspondence theory, but probably the most widespread of these involve doubts about how the relation of correspondence should itself be explicated or clarified—or indeed whether it can be intelligibly explicated at all. It has often been suggested that correspondence must be construed as some sort of complicated structural isomorphism or relation of “picturing” between (a) the components of the proposition, or perhaps of the linguistic expression of the proposition, and (b) the relevant chunk or aspect of reality. And intelligibly defining or specifying a relation of this sort has been argued to be difficult or perhaps impossible. One reason offered for this claim is that the relation would have to be realized by propositions about very widely different sorts of subject matter, such as concrete physical situations (as in the example just discussed), general physical laws, historical facts, facts about mental states, abstract logical and mathematical facts, and perhaps normative or valuational facts. But how, it is asked, could there be one relation of the sort in question that is realized in cases as different as these? How could the very same relation that obtains between the proposition that my car is in the parking lot and the physical configuration described earlier also obtain between the proposition that $2 + 3 = 5$ and the abstract mathematical fact to which it would presumably have to correspond? A second, perhaps even deeper reason often given for thinking that a specification of the correspondence relation is impossible is that to formulate such a specification, we would have to be able to talk about or indicate *both* sides of the relation: both the conceptually formulated, linguistically expressible propositional content and the mind-independent, nonconceptual aspect or chunk of

reality. But, the argument goes, we have no way to get at the latter element except via further conceptual, propositional descriptions, which thus, it is claimed, merely *presuppose* the correspondence relation (assuming for the sake of the argument that the correspondence theory is true) without really helping to explain it. In effect, an attempted account of an instance of the correspondence relation only exhibits a relation between two conceptual, propositional descriptions and not one between such a description and a hunk of independent, nonconceptual reality.

These reasons for doubting whether an intelligible specification of the correspondence relation is possible raise difficult issues, and it would take quite a bit of discussion, more than there is room for here, to get to the bottom of them. Fortunately, however, there is a way of seeing that the problems they raise, though perhaps important in other ways, need not be solved in order to make sense of the correspondence theory of truth itself. The mistake that is made by these reasons and the objection that they support is thinking that the intelligibility of the correspondence theory requires a generally applicable specification of the relation of correspondence in the way that they suppose, at least if such a specification is supposed to be more than utterly straightforward and trivial. Any intelligible proposition, after all, says that reality (in the broadest sense of the term) is a certain way or has certain features that the content of the proposition specifies. And the best way to understand the correspondence theory, following Aristotle's original statement of it,⁵ is to construe it as saying no more than that such a proposition is true if reality is *whatever* way or has *whatever* features the proposition describes it as having. In some cases, the content of a supposed proposition may be less than fully clear or intelligible, but that is a problem for that supposed proposition and not for the correspondence theory. A way of putting this point is to say that the only specification needed as to how reality would have to be to correspond to a particular proposition and so of what correspondence for that particular proposition involves is provided by the propositional content itself and need not be independently specified by the correspondence theory—a point that also allows for very different sort of propositions to describe and so correspond to reality in their own distinctive ways. (There is no room here for a consideration of the various other objections that have been raised against the correspondence theory, though none of these seems to me in the end to have any more force than the one just discussed.)

The belief that the correspondence theory is untenable has also led philosophers to propose a variety of alternative theories or accounts of truth, some of which have also been motivated by related doubts as to whether truth would be knowable or accessible if understood in the way indicated by

the correspondence theory. These views cannot be discussed in any detail here, but a brief enumeration of the most important alternatives and their main problems may help to give you some idea of what they involve:

(1) *The coherence theory of truth.* According to this view, the truth of a believed proposition simply consists in its fitting together coherently with other propositions that are believed, where coherence involves both logical consistency and (usually) other relations of mutual support or explanation. (It is important to understand that this is supposed to be what truth ultimately *amounts to*, not merely—which would be substantially more plausible—a test or criterion for determining what is true.⁶) Since this view seems implicitly to deny the existence of any objects of knowledge beyond beliefs and their propositional contents (for admitting such objects would lead inevitably back to the correspondence theory), it seems to require an idealist metaphysics in which only mental states (and the minds that have them?⁷) genuinely exist.

In addition to the intuitive implausibility of this idealist view, there is the further objection (among others) that it seems possible for there to be many different and incompatible coherent systems of believed propositions, all of the members of which would be true according to the coherence theory—which appears (think about it) to be an absurd result. (This point is sometimes made by suggesting that the propositions reflected in a well-written novel might seemingly satisfy the requirement of coherence, so that the beliefs of someone who accepted all of them would thereby be true according to the coherence theory.)

(2) *The pragmatic theory of truth.* There are a number of different versions of this view, but we will limit ourselves here to the simplest, advanced by the American pragmatist William James,⁸ which holds that the truth is what “works”: that is, that for a believed proposition to be true is for the holding of that belief to lead in general to success in practice. Now there can be little doubt that believing true propositions often leads to success in this way and that believing false propositions often leads to failure: for example, if I have a true rather than a false belief about the location of my car, then my efforts to get to it and drive home are obviously much more likely to succeed.⁹ But is such a belief true *because* it produces success, since producing success is just what truth is (as the pragmatic theory claims)? Or isn’t it exactly the other way around: doesn’t the belief lead to success because it is true (in the correspondence sense)?

(3) *The redundancy or “disappearance” theory of truth.* Some recent philosophers, seeking to avoid the problems that (as they see it) arise from the correspondence theory and these other theories of truth, have suggested that

there is really no need for any philosophical theory of the nature of truth. They point out the necessary equivalence between assertions of the form “P is true,” for some proposition P, and the simple assertion that P, for example between the assertion *that it is true that my car is in the parking lot* and the assertion simply *that my car is in the parking lot*: if one of these claims is true, then the other must be true also, and vice versa (think about it). But this equivalence means, they argue, that the assertion that a particular proposition is true means or says no more than the simple assertion of that proposition, in which case the former can always be replaced by the latter, and any mention of truth thus disappears. Their conclusion is that talk of truth is simply *redundant*—nothing more than a needlessly elaborate way of asserting the propositions in question.

One problem with this view is that there are cases where the claim is made that something is true but where the proposition in question is left unstated (“what Tom said was true”), so that the proposed replacement doesn’t work. A deeper problem is that the equivalence after all works both ways, and thus could at least as reasonably be taken to show that all propositional assertions are implicitly assertions that the proposition in question is true—which would make an understanding of truth essential for even understanding the idea of assertion or belief.

Having briefly canvassed these alternatives, I propose to follow common sense and the main weight of philosophical opinion by assuming that it is the correspondence theory that gives the correct account of truth, and understanding the second condition for knowledge accordingly.¹⁰

The Reason or Justification Condition

The easiest way to understand the need for this third general condition as a part of the concept of knowledge is to consider briefly the suggestion that no such condition (and no other condition beyond the two already discussed) is necessary, that knowledge can be correctly understood merely as true belief. In fact, there are a few kinds of situation where it does seem reasonably natural to say that a person knows something, even though only these two conditions are satisfied. Here the clearest examples are cases where (a) some secret is being hidden from someone, and (b) in which the purpose for the secrecy will be defeated if the person being kept in the dark comes to have a confident true belief about the matter in question, whether or not he or she has any reason or justification at all for this belief. Thus, for example, if I am hiding from Susan (whether seriously or in a game) and she confidently guesses my location and heads in that direction, I might say to myself or to

someone hiding with me “Susan knows where I am hiding”—even if it is just a hunch for which she has no basis at all or even if it results somehow from some sort of mistake or confusion on her part.

But apart from such relatively rare cases in which the truth of the belief in question is in effect all that matters, it seems clear (a point that has been recognized at least since Plato¹¹) that a mere lucky guess or hunch does not suffice for knowledge even though it undeniably may produce a true belief. Really clear illustrations of this point are not easy to find because it is unusual for a person to believe confidently that something is so even when he or she lacks any real basis for the belief. But what is clear is that in a case where the proposition believed happens to be true only by mere luck or accident, a person does not come to know merely by somehow managing to have a sufficiently confident belief. Thus, for example, if a person on a multiple-choice type quiz show has no idea at all about the answer to a particular question and simply hits the right answer by luck, it would be mistaken to ascribe knowledge to them (prior to their being told that the answer was correct) even if they did manage to believe confidently that the choice was correct. (Here again we have an appeal to intuition.) Similarly, a rabid sports fan who is utterly sure that his team will win a certain game even though there is no real evidence or other basis for this claim did not know beforehand that his team would win even if in fact it does. And even in the case discussed earlier, it would be easy to challenge the claim that Susan really *knows* where I am. The right account seems to be rather that in that specific case (and some others) a true belief is just as good as knowledge and can therefore, in what amounts to a kind of exaggeration, be described as such.

What more then is needed for knowledge than a true belief, perhaps a very highly confident one? The answer offered by the traditional conception of knowledge is that one further ingredient is needed: a sufficiently strong *reason* or *justification* for thinking that the claim in question is true. Here the last part of the specification is essential, for there are other sorts of reasons or justification that I might have for holding a belief that would not be of the right kind to yield knowledge. I might believe something out of loyalty to a friend or out of commitment to a religious tradition (also a sort of loyalty) or perhaps even just because it makes me happier to do so, but such beliefs do not thereby constitute knowledge even if they should happen to be true. What is needed for knowledge, according to the traditional conception, is a reason or justification of a sort that is *truth-conducive*: one that increases or enhances (to the appropriate degree—see below) the likelihood that the belief is true. Such a reason or justification is standardly referred to as an *epistemic* reason or as *epistemic* justification.

The most familiar and obvious way to have an epistemic reason for something that I believe is to have *evidence* in favor of the truth of the proposition in question. In the clearest sort of case, evidence consists in further information of some appropriate sort in light of which it becomes *evident* that the proposition is true. Thus, for example, a police detective might have evidence in the form of fingerprints, eyewitness testimony, surveillance photographs, and the like, pointing strongly to the conclusion that a particular person is guilty of the crime he or she is investigating. A scientist might have evidence in the form of instrumental readings and laboratory observations in favor of the truth of a particular scientific theory. And a historian might have evidence in the form of manuscripts and artifacts for the occurrence of a particular historical event.

It is less clear whether the concept of evidence can be extended to encompass all cases in which someone has an epistemic reason or epistemic justification. From an intuitive standpoint, it seems clear that my belief that $2 + 3 = 5$ is epistemically justified, that I have a reason or basis of some sort for thinking that it is true. (I am not merely guessing, nor am I accepting the claim on the basis of authority; rather, I see or grasp directly why the claim is true, indeed why it must be true.) But do I really have evidence that supports the proposition in question? If so, what exactly is it? As we saw in the previous chapter, philosophers have spoken in cases of this kind of *self-evidence*, where this seems to mean that the very content of the proposition in question somehow provides or constitutes evidence for its own truth. We will investigate this idea of self-evidence more fully later on,¹² but it is clear at least that self-evidence does not involve evidence in the most ordinary sense—that is, it does not involve a separate body of information that supports the proposition in question, for otherwise it would not be *self-evidence*.

There are still other sorts of cases of apparent epistemic reasons or justification to which the concept of evidence does not comfortably apply. What about cases of ordinary sensory perception, for example, my present perception of a large green coniferous tree outside my window? Do I have *evidence* for the existence and character of the tree, and if so what might it be? Philosophers have sometimes spoken in such cases of “the evidence of the senses,” but it is far from obvious how this idea should be understood or, here again, that it involves a separate body of supporting information. (Though it is worth noting that philosophers have also sometimes spoken of “sensory information”—can you see anything in such a case that this might refer to?) What about cases of memory? I believe and seem to know that I had Grape-Nuts for breakfast this morning, but do I have *evidence* for this claim when I simply remember it (as opposed to checking the traces left in the

bowl)? And what about my apparent knowledge of my own states of mind, of my “immediate experience”? I believe and seem to know that I am currently thinking about the concept of knowledge, that there is a large patch of dark green in my visual field, that I have an itch in my left elbow, and that I am determined to finish this chapter today, but do I have *evidence* for any of these claims? In this last sort of case, philosophers have also sometimes appealed to the idea of self-evidence. This too will be considered later, but we can see immediately that this is again a rather strained use of the ordinary notion of evidence.

All of these matters will require further discussion later on in this book. For the moment, we can say that the concepts of an epistemic reason or of epistemic justification as they figure in the traditional concept of knowledge are, if not simply identical to the concept of evidence, at least fairly straightforward generalizations of that concept. First, they involve a basis of some sort for thinking that the proposition in question is true or likely to be true, even if not necessarily the sort of separate body of information that the idea of evidence most naturally suggests. Second, on the most standard and obvious interpretation, these concepts also seem to involve the idea that this truth-conducive basis is something that is within the cognitive possession of the person whose belief thereby comes to be justified, that is, that it is something that he or she is aware of in some way that would allow it to be cited as a reason or as giving justification for the belief in question.¹³

Yet a further issue pertaining to the reason or justification condition for knowledge is how *strong* the reason or justification must be, that is, how likely it must make it that the proposition in question is true, for knowledge to result. We have already seen Descartes’s apparent view on this point: the reason must be *conclusive*, must *guarantee* the truth of the proposition in the sense that it is impossible for the proposition to be false, given that reason. Accounts of knowledge that, like the Cartesian account, involve this strong version of the reason or justification condition are sometimes referred to as versions of the “strong conception of knowledge” (or the strong sense of the term “knowledge”).¹⁴

It is fairly easy to see the appeal of the strong version of the reason or justification condition and the strong conception of knowledge that results. If, as suggested earlier, the aim of our cognitive endeavors is truth and our reasons or justification are our means for achieving this goal, then only a reason or justification that satisfies the strong version of the condition allows us to be sure that the goal has in fact been achieved; with anything less than this, success would be to some extent uncertain. Moreover, this interpretation of the third condition for knowledge seems to agree with at least some

of the ways in which we ordinarily use the term “know”: given inconclusive evidence, it is natural for a person to say, at least if pushed, that he or she doesn’t really *know* that the claim in question is true despite having a fairly good reason for believing it.

But the main problem with the strong conception of knowledge is that there seem to be many, many cases that we commonsensically or intuitively regard as cases of knowledge where the strong version of the reason or justification condition is clearly not satisfied. As we have learned from Descartes (even though he himself seems sometimes to lose sight of this lesson in the later stages of the *Meditations*), it is *very* hard to find beliefs for which there is not some *possible* way in which the proposition in question could be false in spite of the reasons or justification for thinking that it is true. Given possibilities like the evil genius, it is doubtful whether any beliefs about the material world outside of our minds or about the past will count as knowledge, according to the strong conception. Indeed, contrary to Descartes, it can even be questioned whether beliefs about our own states of minds will constitute knowledge according to this strong standard: is it really *impossible* (given my evidence or basis, whatever exactly it is) that I could be mistaken about whether I am experiencing a specific shade of color or about how severe a sensation of pain is? Thus if the strong conception is the right account of knowledge, it may well follow that we have virtually no knowledge at all, perhaps nothing beyond the minimal knowledge for each of us of his or her own existence. And this result seems to conflict both with common-sense intuition and with our ordinary usage of the terms “know” and “knowledge.”

It is this sort of objection that has led most recent philosophers to adopt versions of what is sometimes referred to as the “weak conception of knowledge” (or the weak sense of “knowledge”).¹⁵ According to these views, the correct version of the reason or justification condition does not require *conclusive* reasons or justification for there to be knowledge. What is required is instead only *reasonably strong* reasons or justification, strong enough to make it quite *likely* that the proposition in question is true, but not necessarily strong enough to *guarantee* its truth. It is at least fairly plausible to suppose that most or all of the beliefs that we intuitively regard as cases of knowledge do in fact satisfy this less demanding condition.¹⁶

In fact, there is a connection here between the first and the third conditions of any particular version of the traditional conception of knowledge. It seems to be plainly irrational for a person to believe something more strongly than the strength of their reason or justification would warrant (and perhaps also, though less obviously so, to believe it less strongly). Thus if we assume, reasonably enough it would seem, that knowledge in-

volves beliefs that are rationally held, then accepting the weak version of the reason or justification condition is a good reason for also accepting the weaker version of the belief or acceptance condition that was mentioned at the end of the discussion of that condition; whereas one who accepts the strong version of the reason or justification condition has no reason not to accept (and perhaps good reason in favor of accepting) a comparably strong version of the belief or acceptance condition.

One very obvious question to ask about the weak conception is *how* likely the truth of the proposition must be to satisfy this weaker version of the reason or justification condition. If, as seems at least initially reasonable, the level of likelihood can correctly be thought of as something like a level of probability, then just *how* probable must it be in light of the reasons or justification available that the proposition is true in order for it to be adequately justified to count as knowledge? Presumably more than mere 51 percent probability is required, since it seems intuitively wrong to say that a person knows something that is only barely more likely to be true than false—and, of course, obviously wrong to say that something that is less likely to be true than false is known. But how much more is required? Is 80 percent probability adequate or is that still too low? Should it be 90 percent, or 95 percent, or 99 percent, or 99.9 percent? There is no very obvious way of answering this question, and the even more striking fact is that almost none of the advocates of the weak conception of knowledge have ever seriously tried to do so.¹⁷ Even more important, it is simply unclear what sort of basis or rationale there might be for fixing this level of justification in a nonarbitrary way. However problematic the strong conception of justification may be in other ways, its intuitive significance and importance is clear. But nothing like this seems to be true for the weak conception.

This last problem calls into serious question whether any clearly motivated version of the supposed weak conception of knowledge even exists as an alternative to the strong conception. But it will be convenient to defer further discussion of this issue until we have considered a quite different and somewhat surprising problem that has been recently (by philosophical standards) raised in relation to the traditional conception of knowledge.

The Gettier Problem

It is reasonable to say the some version or other of the traditional conception of knowledge was taken for granted, often without very much in the way of detailed specification, by virtually all philosophers seriously concerned with knowledge in the period from the time of Descartes until the middle of the