

5

EPISTEMIC NORMS

1. Recapitulation

We have surveyed existing theories of knowledge and concluded that most are subject to fatal objections. Doxastic theories, both foundationalist and coherence, fail because they cannot accommodate perception and memory. These are cognitive processes that produce beliefs in us, and the beliefs are sometimes justified and sometimes unjustified, but whether they are justified is not just a function of one's other beliefs. It follows that justifiability is a function of more than doxastic states, and hence the true epistemological theory must be a nondoxastic theory. Nondoxastic theories can be internalist or externalist. We have sketched an internalist nondoxastic theory—direct realism—and one of our ultimate purposes in this book is to defend a variety of direct realism. This will be done by arguing against externalist theories. If all externalist theories can be rejected, the only remaining theories are internalist nondoxastic theories, and we take it that direct realism is the most plausible such theory. The premiere externalist theories are all versions of either probabilism or process reliabilism. These theories fail for a variety of reasons specific to them. But it remains possible at this stage that some other form of externalism might succeed, so a more general argument against externalism is required if we are to defend direct realism in this way.

On another front, all of the theories thus far discussed are subject to a common objection. This is that they fail to give illuminating general accounts of epistemic justification. Although they may start with simple and intuitive ideas, when confronted with detailed objections they are forced to complicate those simple ideas and, in the end, they propound complex and convoluted criteria of justifiedness. Even if some such complex criterion were correct, we would be left wondering why such a concept of epistemic justification should be of interest to us. It has been objected that foundations theories and direct realism propose ad hoc lists of epistemic rules whose only defense is that they seem to be required for the justifiedness of those beliefs we antecedently regard as justified. As formulated, those theories give no principled account of epistemic justification from which this medley of rules might be derived. But we have found that much the same objection can be raised to all the other theories we have discussed as well. The final versions of these theories leave us with such complicated criteria that they cannot be regarded as explanations of what epistemic justification is all about.

To sort this out we need a general account of epistemic justification, and it will be the purpose of this chapter to provide such an account.

Once we have a better understanding of epistemic justification it will become possible to dismiss all externalist theories for deep reasons having to do with the general nature of epistemic justification. Basically the same considerations will also necessitate the rejection of a wide variety of internalist theories, including most coherence theories. Ultimately, we aim to offer a simple and explanatory theory of what justification is. The general account of epistemic justification that will be proposed here and in the next chapter has the further virtue that it is a naturalistic account, in the sense that it integrates the concept of epistemic justification into a naturalistic view of human beings as biological machines.

2. Epistemic Norms

What are we asking when we ask whether a belief is justified? What we want to know is whether it is all right to believe it. Justification is a matter of “epistemic permissibility”. It is this normative character of epistemic justification that we want to emphasize. That epistemic justification is a normative notion is not a novel observation. The language of epistemic justification is explicitly normative, and a recurrent theme has been that justification is connected with the “ethics of belief”. This has played a role in the thought of a number of epistemologists. Roderick Chisholm (1977 and chapter one of 1957) has repeatedly stressed the normative character of epistemic terms, several recent philosophers have proposed analyzing epistemic justification in terms of the maximization of epistemic values,¹⁰² and a few philosophers have appealed to the normative character of justification in other ways.¹⁰³ Thus we will think of epistemic justification as being concerned with questions of the form, “When is it permissible (from an epistemological point of view) to believe *P*?” This is the concept of epistemic justification that we are concerned to explore.

Norms are general descriptions of the circumstances under which various kinds of normative judgments are correct. Epistemic norms are norms describing when it is epistemically permissible to hold various beliefs. A belief is justified if and only if it is licensed by correct epistemic norms. We assess the justifiedness of a belief in terms of the cognizer’s reasons for holding it, and our most fundamental epistemic judgments pertain to reasoning (construing reasoning in the broad manner required by direct realism). Thus we can regard epistemic norms as the norms governing “right reasoning”. Epistemic norms are supposed to guide us in reasoning and thereby in forming beliefs. The concept of epistemic

102. See for example Isaac Levi (1967), Keith Lehrer (1974, p. 146ff and 204ff) and (1981, p. 75ff), and Alvin Goldman (1981, pp. 27-52).

103. See for example Hilary Kornblith (1983). See also William Alston (1978), Roderick Firth (1978), John Heil (1983), and Jack Meiland (1980).

justification can therefore be explained by explaining the nature and origin of the epistemic norms that govern our reasoning. We have been calling this “the procedural concept of epistemic justification”. There may be other concepts that can reasonably be labeled “epistemic justification”, but it is the procedural concept that is the focus of the present book and is involved in traditional epistemological problems.

Much of recent epistemology has been concerned with describing the contents of our epistemic norms, but the nature and source of epistemic norms has not received much attention. Epistemologists have commonly supposed that epistemic norms are much like moral norms and that they are used in evaluating reasoning in the same way moral norms are used in evaluating actions. One of the main contentions of this chapter will be that this parallel is not at all exact and that epistemologists have been misled in important ways by supposing the analogy to be better than it is.¹⁰⁴ A proper understanding of epistemic norms will provide us with a radically new perspective on epistemology, and from the point of view of this perspective new light can be shed on a number of central epistemological problems.

3. How Do Epistemic Norms Regulate?

3.1 *Epistemic Normativity*

In order to get a grasp of the nature of epistemic norms, let us begin by asking their purpose. It is important to distinguish between two uses of norms (epistemic or otherwise). On the one hand, there are third-person uses of norms wherein we use the norms to evaluate the behavior of others. Various norms may be appropriate for third-person evaluations, depending upon the purpose we have in making the evaluations. For example, we may want to determine whether a person is a good scientist because we are trying to decide whether to hire her. To be contrasted with third-person uses of norms are first-person uses. First-person uses of norms are, roughly speaking, action-guiding.¹⁰⁵ For example, I might appeal to *Fowler’s Modern English Usage* to decide whether to use ‘that’ or ‘which’ in a sentence. We will call such action-guiding norms “procedural”. Epistemological questions are about rational cognition—about how cognition rationally ought to work—and so are inherently first-person. The traditional epistemologist asks, “How is it possible for me to be justified in my beliefs about the external world, about other minds, about the past, and so on?” These are questions about what to believe. Epistemic norms are the norms in terms of which these questions are to be answered, so these norms are used in a first-person reason-guiding or procedural capacity.

104. See Michael DePaul (1993) for some discussion of this.

105. We can also make “third-person evaluations” of our own past behavior, but that is different from what we are calling “first-person uses” of norms.

3.2 *The Intellectualist Model*

If reasoning is governed by epistemic norms, just how is it governed? There is a model of this regulative process that is often implicit in epistemological thinking, but when we make the model explicit it is *obviously* wrong. This model assimilates the functioning of epistemic norms to the functioning of explicitly articulated norms. For example, naval officers are supposed to “do it by the book”, which means that whenever they are in doubt about what to do in a particular situation they are supposed to consult explicit regulations governing all aspects of their behavior and act accordingly. Explicitly articulated norms are also found in driving manuals, etiquette books, and so on. Without giving the matter much thought, there is a tendency to suppose that all norms work this way, and in particular to suppose that this is the way epistemic norms work. We will call this “the intellectualist model”.¹⁰⁶ It takes little reflection to realize that epistemic norms cannot function in accordance with the intellectualist model. If we had to make an explicit appeal to epistemic norms in order to acquire justified beliefs we would find ourselves in an infinite regress, because to apply explicitly formulated norms we must first acquire justified beliefs about how they apply to this particular case. For example, if we are to reason by making explicit appeal to a norm telling us that it is permissible to move from the belief that something looks red to us to the belief that it is red, we would first have to become justified in believing that that norm is included among our epistemic norms and we would have to become justified in believing that we believe that the object looks red to us. In order to become justified in holding those beliefs, we would have to apply other epistemic norms, and so on *ad infinitum*. Thus it is clear that epistemic norms cannot guide our reasoning in this way.¹⁰⁷

3.3 *Do Epistemic Norms Regulate?*

If the intellectualist model is wrong, then how do epistemic norms govern reasoning? At this point we might raise the possibility that they do not. Perhaps epistemic norms are only of use in third-person evaluations. But it cannot really be true that epistemic norms play *no role at all* in first-person deliberations. We can certainly subject our reasoning to self-criticism. Every philosopher has detected invalid arguments in his or her own reasoning. This might suggest that epistemic norms are

106. Many philosophers appear to adopt the intellectualist model, although it is doubtful that any of them would seriously defend it if challenged. For example, Alvin Goldman (1981) appears to assume such an account of epistemic norms. The intellectualist model pervades Hilary Kornblith’s (1983) discussion. Unfortunately, it is also prominent in Pollock’s (1979) discussion.

107. This point has been made several times. Pollock made it in his (1974), and James van Cleve (1979) made it again. Despite this, we do not think that epistemologists have generally appreciated its significance.

only relevant in a negative way. Our reasoning is innocent until proven guilty. We can use reasoning to criticize reasoning, and hence we can use reasoning in applying epistemic norms to other reasoning, but we cannot be required to reason about norms *before* we can do any reasoning. This would avoid the infinite regress.

As theoretically attractive as the “innocent until proven guilty” picture might be, it cannot be right. It entails the view, already discussed and rejected in chapter three, according to which all beliefs are *prima facie* justified. This view cannot handle the fact that epistemic norms guide the acquisition of beliefs and not just their after-the-fact evaluation. This was illustrated by the observation that even in the perceptual acquisition of beliefs about physical objects, the resulting beliefs are sometimes unjustified. More generally, there are a number of natural processes that lead to belief formation. Among these are such “approved” processes as vision, inductive reasoning, deductive reasoning, and memory, and also some “unapproved” but equally natural processes such as wishful thinking. The latter is just as natural as the former. Recall the example given earlier. My daughter had gone to a football game, the evening had turned cold, and I was worried about whether she took a coat. I found myself thinking, “Oh, I am sure she is wearing a coat”. But then on reflection I decided that I had no reason to believe that—my initial belief was just a matter of wishful thinking. The point here is that wishful thinking is a natural belief-forming process, but we do not accord it the same status as some other belief-forming processes like vision. Although we have a natural tendency to form beliefs by wishful thinking, we also seem to “naturally” know better. This is not just a matter of after-the-fact criticism. We know better than to indulge in wishful thinking at the very time we do it. It seems that *while* we are reasoning we are being guided by epistemic norms that preclude wishful thinking but permit belief formation based upon perception, induction, and so on. This is of more than casual significance, because it might be impossible to rule out wishful thinking by after-the-fact reasoning. This is because the after-the-fact reasoning might include wishful thinking again, and the new wishful thinking could legitimize the earlier wishful thinking. If epistemic norms play no regulative role in our reasoning while it is going on, there is no reason to think they will be able to play a successful corrective role in after-the-fact evaluations of reasoning. In order for the corrective reasoning to be successful it must itself be normatively correct. Epistemic norms must, and apparently do, play a role in guiding our epistemic behavior at the very time it is occurring. But how can they?

Epistemic norms cannot play a merely negative, corrective, role in guiding reasoning, nor can they function in a way that requires us to already make judgments before we can make judgments. What is left? Our perplexity reflects an inadequate understanding of the way procedural norms usually function. The case of making an explicit appeal to norms in order to decide what to do is the exception rather than the rule. You may make reference to a driving manual when you are first learning to

drive a car, but once you learn how to drive a car you do not look things up in the manual anymore. You do not usually give any explicit thought to what to do—you just do it. This does not mean, however, that your behavior is no longer guided by those norms you learned when you first learned to drive. Similarly, when you first learned to ride a bicycle you were told to turn the handlebars to the right when the bicycle leaned to the right. You learned to ride in accordance with that norm, and that norm still governs your bike riding behavior but you no longer have to think about it. The point here is that norms can govern your behavior without your having to think about them. The intellectualist model of the way norms guide behavior is almost always wrong. This point has been insufficiently appreciated. It is of major importance in understanding epistemic norms. Reasoning is more like riding a bicycle than it is like being in the navy.

3.4 *Procedural Knowledge*

What makes it possible for your bike-riding behavior to be governed by norms without your thinking about the norms is that you *know how* to ride a bicycle. This is *procedural knowledge* rather than *declarative knowledge*. Having procedural knowledge of what to do under various circumstances does not involve being able to give a general description of what we should do under those circumstances. This is the familiar observation that knowing how to ride a bicycle does not automatically enable one to write a treatise on bicycle riding. This is true for two different reasons. First, knowing how to ride a bicycle requires us to know what to do in each situation *as it arises*, but it does not require us to be able to say what we should do before the fact. Second, even when a situation has actually arisen, our knowing what to do in that situation need not be propositional knowledge. In the case of knowing that we should turn the handlebars to the right when the bicycle leans to the right, it is plausible to suppose that most bicycle riders do have propositional knowledge of this; but consider knowing how to hit a tennis ball with a tennis racket. We know how to do it—as the situation unfolds, at each instant we know what to do—but even at that instant we cannot give a description of what we should do. Knowing what to do is the same thing as knowing to do it, and that need not involve propositional knowledge.

We can give a rough description of how procedural norms govern behavior in a non-intlectualist manner. When we learn how to do something *X*, we “acquire” a plan of how to do it. That plan might (but need not) start out as explicit propositional knowledge of what to do under various circumstances, but then the plan becomes internalized. Using a computer metaphor, psychologists sometimes talk about procedural knowledge being “compiled-in”. When we subsequently undertake to do *X*, our behavior is automatically channeled into that plan. This is just a fact of psychology. We form habits or conditioned reflexes. Norms for doing *X* constitute a description of this plan for doing *X*. The sense

in which the norms guide our behavior in doing *X* is that the norms describe the way in which, once we have learned how to do *X*, our behavior is automatically channeled in undertaking to do *X*. The norms are not, however, just descriptions of what we do. Rather, they are descriptions of what we *try* to do. Norms can be hard to follow and we follow them with varying degrees of success. Think, for example, of an expert golfer who knows how to swing a golf club. Nevertheless, he does not always get his stroke right. It is noteworthy, and it will be important later, that when he does not get his stroke right he is often able to tell that by something akin to introspection. When he does it wrong it “feels wrong”. The ability to tell in this way whether one is doing something right is particularly important for those skills governing performances (like golf swings) that take place over more than just an instant of time, because it enables us to correct or fine tune our performance as we go along.

The internalization of norms results in our having “automatic” procedural knowledge that enables us to do something without having to think about how to do it. It is this process that we are calling “being guided by the norm without having to think about the norm”. This may be a slightly misleading way of talking, because it suggests that somewhere in our heads there is a mental representation of the norm and that mental representation is doing the guiding. Perhaps it would be less misleading to say that our behavior is being guided by our procedural knowledge and the way in which it is being guided is described by the norm. What is important is that this is a particular way of being guided. It involves non-intellectual psychological mechanisms that both guide and correct (or fine tune) our behavior.

3.4.1 The Competence/Performance Distinction

The distinction between knowing how to do something and actually doing it is the same as the competence/performance distinction in linguistics. When linguists study a language, they try to discover what the rules are that determine which utterances are grammatical. But it is a contingent matter what the rules are that govern any given language, and those rules may change over time as the language evolves. The language is determined by the way the speakers of the language use it. However, eliciting the rules of grammaticality for a language is not the same thing as simply describing how the users of the language talk. Linguists observe that many, perhaps most, of our utterances are ungrammatical. Our speech is populated with “Ahh”s and “Umm”s, we leave sentences unfinished, and commit a variety of other grammatical infractions. But we know better. When linguists study language, they are not interested in what we do when we do it wrong. Linguistic theories are about what we do when we do it right. However, right and wrong in this case are not determined by some metaphysically necessary standard. They are determined by how people talk. To prevent the account from going around in a circle, Noam Chomsky (1965) proposed

that what people know when they know how to speak a language takes the form of procedural knowledge. They *know how* to speak the language. As in the case of any procedural knowledge, people can have internalized rules for how to speak the language but violate them. The objective of the linguist is to describe the internalized rules that comprise the speakers' procedural knowledge. Chomsky referred to this as a *competence theory* of language, and contrasted it with a *performance theory*, which would be a theory of what people actually say, errors and all. A performance theory might be the kind of theory that a psychologist would produce, but linguists seek a competence theory.

3.4.2 Normative Language

The use of normative language in formulating procedural norms is pervasive. Norms can be described as knowing what we *should* do under particular circumstances. The point of using normative language to describe internalized norms is to contrast what the norms tell us to do with what we *do*. The simple fact of the matter is that even when we know how to do something (e.g., swing a golf club) we do not always succeed in following our norms. This use of 'should' in describing procedural knowledge is interesting. Moral philosophers have talked about different senses of 'should', distinguishing particularly between moral uses of 'should' and goal-directed uses of 'should'. An example of the latter is "If you want the knife to be sharp then you should sharpen it on the whetstone". But the use of 'should' in "In riding a bicycle, when the bicycle leans to the right you should turn the handlebars to the right" is of neither of these varieties. It is perhaps more like the goal-directed kind of 'should', but we are not saying that that is what you should do to achieve the goal of riding a bicycle. Rather, that is part of what is involved *in* riding a bicycle—that is *how* to ride a bicycle.

Note that a similar use of normative language occurs in formulating rules of grammar. We are informed that under certain circumstances we should say 'that' rather than 'which'. The 'should' here is not a moral or prudential should. It is that kind of 'should' that we use in describing procedural knowledge. Because it is natural to use normative language in describing procedural knowledge, it is equally natural to say that in acquiring procedural knowledge, what we learn are norms for how we should do something.

3.4.3 Epistemic Norms Are Procedural Norms

So far we have been describing procedural norms in general. Now let us apply these insights to epistemic norms. We know how to reason, or more generally, how to cognize. That means that under various circumstances we know what to do in cognizing. This can be described equivalently by saying that we know what we should do. Our epistemic norms are just the norms that describe this procedural knowledge, and rational cognition is cognition in compliance with the norms. The way

epistemic norms can guide our cognition without our having to think about them is no longer mysterious. They describe an internalized pattern of behavior that we automatically follow in epistemic cognition, in the same way we automatically follow a pattern in bicycle riding. Epistemic norms are the internalized norms that govern our epistemic cognition. Once we realize that they are just one more manifestation of the general phenomenon of automatic behavior governed by internalized norms, epistemic norms should no longer seem puzzling. The mystery surrounding epistemic norms evaporates once we recognize that the governing process is a general one and its application to epistemic norms and epistemic cognition is not much different from its application to any other kind of procedural norms. Of course, unlike most norms our epistemic norms may be innate, in which case there is no process of internalization that is required to make them available for use in guiding our cognition.

There has been a great deal of recent work in psychology concerning human irrationality. Psychologists have shown that in certain kinds of epistemic situations people have an almost overpowering tendency to reason incorrectly.¹⁰⁸ It might be tempting to conclude from this that, contrary to what we are claiming, people do not know how to reason.¹⁰⁹ The short way with this charge is to note that if we did not know how to reason correctly in these cases, we would be unable to discover that people reason incorrectly. To say that we know how to reason is to invoke a competence/performance distinction. It is no way precludes our making mistakes. It does not even preclude our almost always making mistakes in specific kinds of reasoning. All it requires is that we can, in principle, discover the errors of our ways and correct them.¹¹⁰

4. The Refutation of Externalism

We have described how our epistemic norms work. But this is not yet to say anything about which epistemic norms are correct. An epistemological theory must answer two different questions. First, it must describe the correct epistemic norms. Second, it must tell us what makes them correct. The first question concerns the content of epistemic norms,

108. Much of the central psychological material can be found in Daniel Kahneman, Paul Slovic, and Amos Tversky (1982), and R. E. Nisbett and L. Ross (1980). For an overview, see Massimo Piatelli-Palmarini (1994).

109. Murray Clarke (1990) criticizes our view on these grounds.

110. This is pretty much the same as the assessment of the irrationality literature offered by Jonathan Cohen (1981). See also the critique in Alvin Goldman (1986). Edward Stein (1995) offers a thorough and illuminating discussion of the philosophical import of psychological research on reasoning. It should be noted here that not all psychological research is pessimistic about the performance of human reasoning. See Gerd Gigerenzer (1991) and (1996).

and the second question concerns their justification. By distinguishing between these questions we can see the internalism/externalism distinction in a new light. A belief is justified if and only if it is held in compliance with correct epistemic norms. Externalism is the view that the justifiedness of a belief is a function in part of external considerations. Thus if externalism is right, external considerations must play a role in determining whether a belief is held in compliance with correct epistemic norms. This could arise in either of two ways. On the one hand, external considerations could enter into the formulation of correct epistemic norms. On the other hand, it might be granted that epistemic norms can only appeal to internal considerations, but it might be insisted that external considerations are relevant to determining which set of internalist norms is correct. Thus we are led to a distinction between two kinds of externalism. *Belief externalism* insists that correct epistemic norms must be formulated in terms of external considerations. A typical example of such a proposed norm might be "It is permissible to hold a belief if it is generated by a reliable cognitive process." In contrast to this, *norm externalism* acknowledges that the content of our epistemic norms must be internalist, but employs external considerations in the selection of the norms themselves. The distinction between belief and norm externalism is analogous to the distinction between act and rule utilitarianism. Externalism (simpliciter) is the disjunction of belief externalism and norm externalism. In the last chapter, we were concerned with dealing with proposals of the first sort. At the end of that discussion, we noted that there was another sense of externalism that we would eventually have to grapple with, namely norm externalism. A number of philosophers who are normally considered externalists appear to vacillate between belief externalism and norm externalism.¹¹¹ The difference between these two varieties of externalism will prove important. In the end, both must be rejected, but they are subject to different difficulties.¹¹²

According to internalism, the justifiedness of a belief is a function exclusively of internal considerations, so internalism implies the denial of both belief and norm externalism. That is, the internalist maintains that epistemic norms must be formulated in terms of relations between beliefs or between beliefs and nondoxastic internal states (e.g., perceptual states), and she denies that these norms are subject to evaluation in terms of external considerations. Typically, the internalist has held that

111. Alvin Plantinga (1993a) is guilty of this. He criticizes the view that we are defending in this book by assuming that some version of norm externalism can be made to work (see especially pp. 171-6). He fails to recognize that norm externalism would have to be defended separately from his endorsement of the Theory of Proper Functions, which is a kind of belief externalism.

112. Alvin Goldman (1981) seems to be one of the few externalists who is clear on this distinction. He distinguishes between two senses of 'epistemic justification' and adopts belief externalism with regard to one and norm externalism with regard to the other.

whatever our *actual* epistemic norms are, they are necessarily correct and not subject to criticism on any grounds (externalist or otherwise). Of course, this is precisely where internalists disagree with norm externalists. Let us turn then to a reconsideration of both forms of externalism in the light of our new understanding of epistemic norms.

4.1 *Belief Externalism*

Now that we understand how epistemic norms work in guiding our epistemic cognition, it is easy to see that they must be internalist norms. This is because when we learn how to do something we acquire a set of norms for doing it and these norms are internalized in a way enabling our cognitive system to follow them in an automatic way without our having to think about them. This has implications for the content of our norms. For example, we have been describing one of our bike-riding norms as telling us that if the bicycle leans to the right then we should turn the handlebars to the right, but that is not really what we learn when we learn to ride a bicycle. The automatic processing systems implemented in our neurology do not have access to whether the bicycle is leaning to the right. What they do have access to are things like our *thinking* that the bicycle is leaning to the right, and certain balance sensations emanating from our inner ear. What we learn is (roughly) to turn the handlebars to the right if we either experience those balance sensations or think on some other basis that the bicycle is leaning to the right. The circumstance-types to which our norms appeal in telling us to do something in circumstances of those types must be directly accessible to our system of cognitive processing. The sense in which they must be directly accessible is that our cognitive system must be able to access them without our first having to make a *judgment* about whether we are in circumstances of that type. We must have non-epistemic access.¹¹³

This general observation about procedural norms has immediate implications for the nature of our epistemic norms. It implies, for example, that epistemic norms cannot appeal to external considerations of reliability. This is because such norms could not be internalized. Like *the bicycle's leaning to the right*, considerations of reliability are not directly accessible to our automatic processing systems. There is in principle no way that we can learn to make inferences of various kinds only if they are *in fact* reliable. Of course, we could learn to make certain inferences only if we *think* they are reliable, but that would be an internalist norm appealing

113. It might be insisted that this is at least sometimes a misleading way of talking—if our norms for doing *X* tell us to do *Y* whenever we *think* it is the case that *C*, we might better describe our norms as telling us to do *Y* when it *is* the case that *C*. We do not care if one chooses to talk that way, but it must be realized that it has the consequence that although the reformulated norm says to do *Y* when it is the case that *C*, knowing how to do *X* will really only result in our doing *Y* when we *think* it is the case that *C*. This will be important. (And, of course, norms appealing to internal states other than beliefs could not be reformulated in this manner anyway.)

to *thoughts* about reliability rather than an externalist norm appealing to reliability itself.¹¹⁴ Similar observations apply to any externalist norms. Consequently, it is in principle impossible for us to employ externalist norms. We take this to be a conclusive refutation of belief externalism.

We introduced the internalism/externalism distinction by saying that internalist theories make justifiedness a function exclusively of the believer's internal states, where internal states are those that are "directly accessible" to the believer. The notion of direct accessibility was purposely left vague, but it can now be clarified. We propose to define internal states as those states that are directly accessible to the cognitive mechanisms that direct our epistemic cognition. The sense in which they are *directly* accessible is that access to them does not require us first to have beliefs about them. This definition makes the internalist/externalist distinction precise in a way that agrees at least approximately with the way it has generally been used, although it is impossible to make it agree with everything everyone has said about it because philosophers have drawn the distinction in different ways. It is especially noteworthy, however, that our access constraint is considerably more liberal than the reflective access required by some internalist theories. Reflective access seems too restrictive in light of the way that norms automatically govern cognition. The internalism/externalism distinction will be discussed further when we reflect on the status of naturalism in epistemology.

We have characterized internalist theories in terms of direct accessibility, but we have not said anything in a general way about which properties and relations are directly accessible. It seems clear that directly accessible properties must be in some sense "psychological", but it is doubtful that we can say much more than that from the comfort of our armchairs. That is an empirical question to be answered by psychologists. Despite the fact that we do not have a general characterization of direct accessibility, it is perfectly clear in many specific cases that particular properties to which philosophers have appealed are not directly accessible. In light of this, the preceding refutation of belief externalism can be applied to a remarkably broad spectrum of theories, and it seems to us to constitute an absolutely conclusive refutation of those theories. We have indicated how it applies to theories formulating epistemic norms in terms of reliability. It applies in the same way to probabilist theories. For example, we saw that many probabilists endorse the *simple rule*:

A belief is epistemically permissible if and only if what is believed is sufficiently probable.

114. It would also be a wholly implausible theory. We do not invariably have beliefs about the reliability of our inferences whenever we make them, and if norms *requiring* us to have such beliefs also require those beliefs to be justified then they lead to an infinite regress.

If the simple rule is to provide us with a procedural norm then the probability of a belief must be a directly accessible property of it. No objective probability can have that property. Thus it is impossible to use the simple rule, interpreted in terms of objective probabilities, as a procedural norm. This objection could be circumvented by replacing the simple rule with its “doxastic counterpart”:

A belief is epistemically permissible if and only if the epistemic agent believes it to be highly probable.

But this rule formulates an internalist norm (albeit, an implausible one¹¹⁵). It might be supposed that we could breath life back into the simple rule by interpreting it in terms of subjective probability. Here we must be careful to distinguish between subjective probability as actual degree of belief and subjective probability as rational degree of belief. Interpreted in terms of actual degrees of belief, the simple rule would amount to the claim that a belief is justified if and only if it is firmly held, which is an internalist norm, but a preposterous one. To get a plausible norm, we must interpret subjective probability as rational degrees of belief. Rational degree of belief is the unique degree of belief one rationally ought to have in a proposition given one’s overall doxastic state, and this is to be understood in terms of prudentially rational betting behavior. As we have indicated, we have serious doubts about the intelligibility of this notion. But even if we waive that objection, ascertaining what this unique rational degree of belief should be is immensely difficult. It seems extremely unlikely that the rational degree of belief one ought to have in a proposition is a directly accessible property of it. If it is not then this version of the simple rule also succumbs to our general objection to belief externalism.

Other epistemological theories succumb to this objection to belief externalism. For example, Keith Lehrer’s coherence theory is an internalist theory, but it was pointed out in the last chapter that an externalist theory can be modeled on it. According to this externalist theory, a person is justified in believing a proposition if and only if that proposition is more probable than each proposition competing with it. But a proposition’s being more probable than any of its competitors is not a directly accessible property of it, and hence the objective version of Lehrer’s theory becomes incapable of supplying us with a procedural norm.

The net cast by this objection catches some internalist theories as well. For instance, a holistic coherence theory adopts a holistic view of reasons according to which a belief is licensed if it is suitably related to the set of *all* the beliefs one holds. A holistic coherence theory requires a

¹¹⁵ We do not ordinarily have any beliefs at all about the probabilities of what we believe. Furthermore, even if we did they would presumably not render our beliefs justified unless the probability beliefs were themselves justified, so we would be threatened by an infinite regress.

relationship between a justified belief and the set of all the beliefs one holds, but that will not normally be a directly accessible property of the justified belief, and hence although the norm proposed by the holistic theory will be an internalist norm, it will not be internalizable. Thus it cannot serve as a procedural norm.

The general point emerging from all this is that the norms proposed by many traditional theories cannot be reason-guiding. Accordingly, they cannot serve as epistemic norms. No non-internalist theory can provide us with epistemic norms that we could actually use. Correct epistemic norms must be internalist. On the other hand, we have also seen that epistemic norms must be able to appeal to more than the cognizer's doxastic state. They must also be able to appeal to his perceptual and memory states. Thus the correct epistemological theory must endorse some kind of nondoxastic internalist norm.

Is there any way to salvage belief externalism in the face of the objection that it cannot give reasonable accounts of first-person reason-guiding epistemic norms? The possibility remains that belief externalism might provide norms for third-person evaluations. We think it is noteworthy in this connection that externalists tend to take a third-person point of view in discussing epistemology. If externalist norms played a role in third-person evaluations, we would then have both externalist and internalist norms that could be applied to individual beliefs and they might conflict. What would this show? It would not show anything—they would just be different norms evaluating the same belief from different points of view. We can imagine a persistent externalist insisting, "Well, if the two sets of norms conflict, which way should we reason—which set of norms should we follow?" But that question does not make any sense. Asking what we should do is asking for a normative judgment, and before we can answer the question we must inquire to what norms the 'should' is appealing. To make this clearer consider an analogous case. We can evaluate beliefs from both an epistemic point of view and a prudential point of view. Recall Helen who has good reason for believing that her father is Jack the Ripper. Suppose that if she believed that, it would be psychologically crushing. Then we might say that, epistemically, she should believe it, but prudentially she should not. If one then insists upon asking, "Well, should she believe it or not?", the proper response is, "In what sense of 'should'—epistemic or prudential?" Similarly, if externalist and internalist norms conflict and one asks, "Which way should we reason?", the proper response is to ask to which set of norms the 'should' is appealing. The point is that different norms serve different purposes, and when they conflict that does not show that there is something wrong with one of the sets of norms—it just shows that the different norms are doing different jobs. The job of internalist norms is reason-guiding, and as such they are the norms traditionally sought in epistemology. Externalist norms (if any sense can be made of them) may also have a point, but they cannot be used to solve traditional epistemological problems pertaining to epistemic justification.

4.2 *Reconsidering the Doxastic Assumption*

The endorsement of nondoxastic norms amounts to the rejection of the doxastic assumption, but that has often seemed puzzling. How is it possible for nondoxastic states to justify beliefs when we are not aware that we are in them? Recall the argument given for the doxastic assumption in chapter one. Procedural epistemic justification is supposed to be concerned with what to believe. But in deciding what to believe, we can only take account of something insofar as we have a belief about it. Thus only beliefs can be relevant to what we are justified in believing. We are now in a position to see what is wrong with that argument, and accordingly to understand how nondoxastic norms are possible. First notice that this argument for the doxastic assumption could not possibly be right, because it is self-defeating. If this argument were right, we could only take account of our beliefs insofar as we have beliefs about our beliefs, and then an infinite regress would loom. There has to be something about beliefs that makes them the sort of thing we can take account of without having beliefs about them. What could this be?

What is it to take account of something in the course of cognition? It is to use it in our cognitive deliberations. We can take account of anything by having a belief about it, but cognition has to start somewhere, with things that we don't have beliefs about. Obviously, it can start with beliefs. The reason it can start with beliefs is that they are internal states, and cognition is an internal process that can access internal states directly. Cognition works by noting that we have certain beliefs and using that to trigger the formation of further beliefs. However, it is *cognition* that must note that we have certain beliefs—*we* do not have to note it ourselves. The sense in which cognition notes it is metaphorical—it is the same as the sense in which a computer program accessing a database might be described as noting that some particular item is contained in it.

We have seen that epistemologists have a lamentable tendency to over-intellectualize cognition. Human beings are *cognitive machines*. We are unusual machines in that our machinery can turn upon itself and enable us to direct many of our own internal operations. Many of these operations, like reasoning, can proceed mechanically, without any deliberate direction or intervention from us, but when we take a mind to we can directly affect their course. For example, we can, at least to some extent, decide what to think about, decide not to pursue a certain line of investigation, and to pursue another one instead. There must, however, be a limit to the extent to which we are *required* to do this. After all, the processes by which we do it are a subspecies of the very processes in which we are intervening. If we had to explicitly direct all of our cognitive processes, we would also have to direct the ones involved in doing the directing, and we would again have an infinite regress.

The significance of this is that we don't *have* to think about our reasoning in order to reason. It is important, for various reasons, that we *can* think about it when the need arises, but we don't have to and don't usually do

it. Thus reasoning can proceed by moving from beliefs to beliefs without our thinking about either reasoning or the beliefs. By virtue of doing the reasoning we are thinking about whatever the beliefs are about, not about the beliefs themselves. This explains the sense in which cognition can take account of our having certain beliefs without our having to have beliefs to the effect that we have those beliefs. But note that it explains much more. In precisely the same sense, cognition can take account of other internal states, for example, percepts, without our having to have beliefs to the effect that we are in those states. Thus there is no reason why cognition cannot move directly from percepts to beliefs about the physical objects putatively represented by the percepts.

Cognition can make use of any states to which it has direct access, but those are just the internal states. So cognition can make use of any internal states without our having beliefs about those states, and correspondingly our epistemic norms can appeal to any internal states—not just beliefs. Such nondoxastic norms only seemed puzzling because we were implicitly assuming the intellectualist model of the way epistemic norms regulate belief. Given the way epistemic norms actually operate, all that is required is that the input states be directly accessible. Belief states are directly accessible, but so are a variety of nondoxastic states like perceptual states and memory states. Thus there is no reason why epistemic norms cannot appeal to those states, and the rejection of the doxastic assumption and the move to direct realism ceases to be puzzling.

4.3 *Norm Externalism*

Recall that there are two kinds of externalism. Belief externalism advocates the adoption of externalist norms. We regard belief externalism as having been decisively refuted by the preceding considerations. Norm externalism, on the other hand, acknowledges that we must employ internalist norms in our reasoning, but proposes that alternative sets of internalist norms should be evaluated in terms of external considerations. For example, it may be alleged that one set of internalist norms is better than another if the first is more reliable in producing true beliefs.¹¹⁶

Internalist theories make the justifiability of a belief a function of the internal states of the believer, in the sense that if we vary anything but his internal states the justifiability of the belief does not vary. Thus the only properties of and relations between internal states to which internalist norms can appeal are those that cannot be varied without varying the internal states themselves. In other words, it must be necessarily true that if we are in those states then they have those properties and stand in those relations to one another. In short, they are “logical” properties of and “logical” relations between internal states. For instance, if S_1 is the state of believing $(P \ \& \ Q)$ and S_2 is the state of believing P , then S_1 and S_2 are necessarily related by the fact that being in S_1 involves believing a

116. As Alvin Goldman does in his (1986).

conjunction whose first conjunct is believed if one is in state S_2 . Thus we can characterize internalist theories as those proposing epistemic norms that appeal only to logical properties of and logical relations between internal states of the believer.

So, both internalism and norm externalism endorse internalist norms, but they differ in that, by definition, the internalist maintains that our epistemic norms are not subject to criticism on externalist grounds. It is hard to see how they could be subject to criticism on internalist grounds, so the internalist has typically assumed that our epistemic norms are immune from criticism—whatever our actual epistemic norms are, they are the correct epistemic norms. That, however, seems odd. On the surface, it seems it must be at least logically possible for two people to employ different epistemic norms. They could then hold the same belief under the same circumstances and on the basis of the same evidence and yet the first could be conforming to his norms and the second not conforming to his. If a person's epistemic norms are always beyond criticism, it would follow that the first person is justified in his beliefs and the second is not, despite the fact that their beliefs are based upon the same evidence. That would at least be peculiar. Because it seems that it must be possible for different people to employ different epistemic norms, this makes a strong *prima facie* case for norm externalism.

The *prima facie* case for norm externalism is bolstered when we notice that procedural norms are not generally immune to criticism. Typically, procedural norms tell us how to do one thing *by* doing something else.¹¹⁷ For example, knowing how to ride a bicycle consists of knowing what more basic actions to perform—leg movements, arm movements, and the like—by doing which we ride the bicycle. An action that is performed by doing something else is a *nonbasic* action. Norms describing how to perform nonbasic actions can be subject to external evaluation. There may be more than one way to perform the nonbasic action, and some ways may be better (more efficient, more reliable, and so on) than others. If I know how to do it in one way and you know how to do it in another way, you know how to do it better than I if the norms governing your behavior are better than the norms governing mine. For example, we may both know how to hit the target with a bow and arrow, but you may know how to do it more reliably than I.¹¹⁸ It thus becomes an empirical question whether acting in accordance with a proposed norm will constitute your doing what you want to be doing and whether another norm might not be better.

Reasoning is not, strictly speaking, an action, but it is something we do, and we do it by doing other simpler things. We reason by adopting new beliefs and rejecting old beliefs under a variety of circumstances.

117. The *by*-relation is what Alvin Goldman (1976a) calls *level-generation*.

118. Alternatively, we may have the same norms but your physical skills make you better able to conform to them.

Our norms for reasoning tell us when it is permissible or impermissible to do this. It seems that the norms we actually employ should be subject to external criticism just like any other norms. The norm externalist proposes that we should scrutinize them and possibly replace them by other norms. Because of the direct accessibility problem, we cannot replace them by norms making explicit appeal to reliability, but what we might discover is that (1) under certain circumstances inferences licensed by our natural norms are unreliable, and (2) under certain circumstances inferences not licensed by our natural norms are highly reliable. The norm externalist proposes that we should then alter our epistemic norms, adopting new internalist norms allowing us to make the inferences described under (2) and prohibiting those described under (1).

Some care is required here. We must distinguish between two construals of the norm externalist's proposal. He might be telling us that when we *discover* old reasoning patterns to be unreliable or new reasoning patterns to be reliable then we should alter our norms and our reasoning accordingly. Alternatively, he might be telling us that if old patterns simply *are* unreliable and new patterns *are* reliable, independently of our knowing or believing that they are, then we should alter our reasoning. The first construal seems like an eminently reasonable proposal, and it is one that has been made explicitly by various externalists. For example, in discussing how reliabilist considerations bear on reasoning, Goldman (1981) writes:

At the start a creature forms beliefs from automatic, preprogrammed doxastic processes. ... Once the creature distinguishes between more and less reliable belief-forming processes, it has taken the first step toward doxastic appraisal. ... The creature can also begin doxastic self-criticism, in which it proposes *regulative* principles to itself (p. 47).

But this involves a fundamental misconception. Our epistemic norms are not subject to criticism in this way. Particular instances of reasoning are subject to such criticism, and the criticism can dictate changes in that reasoning, but this does not lead to changes in our epistemic norms. This is because unlike other norms, our epistemic norms already accommodate criticism based on reliability. The point is twofold. First, discovering that certain kinds of inferences are unreliable under certain circumstances constitutes a defeater for those inferences and hence makes us unjustified in reasoning in that way, and this is entirely in accordance with our natural unmodified epistemic norms. For example, we discover that color vision is unreliable in dim lighting, and once we discover this we should cease to judge colors on that basis under those circumstances. But this does not require an alteration of our epistemic norms, because color vision only provides us with defeasible reasons for color judgments, and our discovery of unreliability constitutes a defeater for those reasons. This is entirely in accordance with the norms we already have. Second, discovering that some new inferences are reliable under certain circum-

stances provides us with justification for making those inferences under those circumstances, but this is licensed by the norms we already have. That is precisely what induction is all about. For example, I might discover that I am clairvoyant and certain kinds of “visions” provide reliable indications of what is about to happen. Once I make this discovery it becomes reasonable for me to base beliefs about the future on such visions. Again, this is entirely in accordance with the norms we already have and does not require us to alter those norms in any way. The general point is that the kinds of reliability considerations to which the norm externalist appeals can lead us to reason differently (refrain from some old inferences and make some new inferences), but this does not lead to any change in our epistemic norms. Our actual epistemic norms are self-correcting in that they involve a kind of built-in feedback having the result that the sort of external criticism that could lead to the modification of other procedural norms does not necessitate any modification of epistemic norms.

We have had several externalists respond to this objection by protesting that they do not see the point of distinguishing between considerations of reliability leading us to alter our reasoning and those considerations leading us to alter our norms. But if all the externalist means is that considerations of reliability can lead us to alter our reasoning, then he is not disagreeing with anyone. In particular, he is not disagreeing with paradigmatic internalists like Chisholm. Norm externalism becomes nothing but a pretentious statement of a platitude.

The alternative construal of norm externalism takes it to be telling us that if old patterns of reasoning are unreliable and new patterns are reliable, then regardless of whether we *know* these facts about reliability, we should not reason in accordance with the old patterns and we should reason in accordance with the new patterns. What could the point of this claim be? It cannot be taken as a recommendation about how to reason, because it is not a recommendation anyone could follow. We can only alter our reasoning in response to facts about reliability if we are apprised of those facts. However, normative judgments do not always have the force of recommendations. That is, they are not always intended to be action-guiding. This is connected with the distinction that is often made in ethics between subjective and objective senses of ‘should’. To say that a person subjectively should do *X* is to say, roughly, that given what he believes (perhaps falsely) to be the case he has an obligation to do *X*. To say that he objectively should do *X* is to say, roughly, that if he were apprised of all the relevant facts then he would have an obligation to do *X*. Judgments about what a person subjectively should do can serve as recommendations, but judgments about what a person objectively should do can only serve as external evaluations having some purpose other than guiding behavior.¹¹⁹ The subjective/objective distinction can

119. They may serve as recommendations in an indirect fashion by conveying to a

be regarded as a distinction between evaluating the person and evaluating her act. The subjective sense of 'should' has to do with moral responsibility, while the objective sense has to do with what act might best have been performed.

We can draw a similar subjective/objective distinction in epistemology. The epistemic analogue of moral responsibility is epistemic justification. A person is being "epistemically responsible" just in case her beliefs are justified. In other words, epistemic justification corresponds to *subjective* moral obligation. What determines whether a belief is justified is what else the epistemic agent *believes* about the world (and what other directly accessible states she is in)—not what is in fact true about the world. This seems to show that whatever considerations of de facto reliability may bear upon, it is not epistemic justification. They must instead bear upon the epistemic analogue of objective obligation. What is that analogue? There is one clear analogue—objective epistemic justification is a matter of what you should believe if you were apprised of all the relevant truths. But what you should believe if you were apprised of all the relevant truths is just *all the truths*. In other words, the epistemic analogue of objective justification is *truth*. There is nothing here to give solace to a norm externalist.

Goldman (1981) draws a somewhat different distinction between two senses of 'justified' in epistemology. He distinguishes between "theoretical" evaluations of reasoning and "regulative" evaluations (the latter being reason-guiding). He suggests that the theoretical sense of justification is the sense required for knowledge and that it is to be distinguished from the reason-guiding sense. He suggests further that his reliabilist theory concerns the theoretical sense. The proposal is that it is knowledge that provides the point of a norm externalist's evaluation of epistemic norms in terms of considerations of reliability unknown to the epistemic agent. We do not believe that, but even if it were true it would not affect our overall point. The sense of epistemic justification with which we are concerned in this book is the reason-guiding or procedural sense, and if it is acknowledged that norm externalism bears only upon another sense of justification then our main point has been conceded.

To summarize the discussion of externalism, one can be an externalist by being either a belief externalist or a norm externalist. These exhaust the ways in which externalist considerations might be brought to bear on our epistemic norms. The belief externalist tries to formulate epistemic norms directly in terms of externalist considerations, but it is impossible to construct procedural norms in this way. The norm externalist proposes instead to recommend changes in procedural norms on the basis of considerations of reliability. Norm externalism initially appeared compelling because it provided a way to preserve the internal and procedural nature of epistemic norms while still allowing for external assessment

person that there are relevant facts of which he is not apprised.

of those norms. Combined with internalism's apparent inability to make sense of the comparative evaluation of norms, it seemed that norm externalism was very promising. Unfortunately, norm externalism fails on its two most plausible construals. It either fails to provide anything that reasoning within the framework of internalism does not, or it reduces justification to truth. So, norm externalism must be rejected.

As far as we can see, externalism has nothing to contribute to the solution to traditional epistemological problems. Justified beliefs are those resulting from normatively correct reasoning. Consequently, any evaluation of the justifiedness of a belief must be reason-guiding and hence must be beyond the pale of externalism.

4.4 *Epistemological Relativism and the Individuation of Concepts*

The apparent failure of norm externalism leaves us with a puzzling problem. Internalists have typically assumed that whatever epistemic norms we actually employ are automatically correct. But that seems hard to reconcile with the seemingly obvious fact that it is at least logically possible for different people to employ different norms. Surely, if Smith and Jones believe *P* for the same reasons, they are either both justified or both unjustified. There is no room for their justification to be relative to idiosyncratic features of their psychology resulting in their employing different epistemic norms. This seems to imply that there is just one set of correct epistemic norms, and the norms a person actually employs may fail to be correct. This conclusion would be obvious if it were not for the fact that there is no apparent basis for criticizing a person's norms. That is precisely what norm externalism tries unsuccessfully to do. The reliabilist considerations to which the norm externalist appeals are the only plausible candidates for considerations of use in criticizing and correcting epistemic norms, and we have seen that our epistemic norms cannot be corrected in this way. Of course, I might criticize Jones' norms simply because they disagree with mine, but he could equally criticize mine because they disagree with his. Are we committed to a thoroughgoing epistemological relativism then? That is at least unpalatable.

4.4.1 Theories of Individuation

The solution to the problem of relativism can be found by turning to a different problem. This is the problem of how concepts are individuated. If it could be shown that people who employ different norms are also necessarily employing different concepts in their reasoning, the troubling possibility of relativism would be dispatched. This is because epistemological relativism maintains that people are using the *same* concepts to reason according to different norms. We aim to show that people using different norms are employing different concepts, and this requires a substantial detour through a theory of the individuation of concepts. The detour will, however, yield a considerable payoff. We will secure a way to avoid relativism about epistemic norms even in the wake of the failure of norm externalism.

To understand the nature of the problem of concept individuation, first consider an analogous problem—that of object individuation. A theory of object individuation is a theory of what makes a physical object the object that it is, and by virtue of what two different objects are different. The historically most popular theory of object individuation proposes to individuate objects in terms of spatio-temporal continuity. On this account, object x and object y are the same object just in case they occupy the same space at the same time. Whatever you may think about the truth of this claim, it is a substantive theory and it is attempting to tell us something nontrivial about physical objects.

4.4.2 Truth Conditions

Theories of the individuation of concepts are similar. They attempt to tell us when concept A is the same concept as concept B . The standard theory takes concepts to be individuated by their truth conditions. The claim of this theory is that what makes a concept the concept that it is are the conditions that must be satisfied for something to exemplify that concept. These conditions comprise its truth conditions. The precise content of the truth condition theory of concepts deserves closer inspection than it usually receives. There is one sense in which the truth condition theory of concepts is correct but also completely trivial and uninteresting. The truth condition of the concept *red* is the condition of *being red*, and the truth condition of the concept *blue* is the condition of *being blue*. The following is undeniable:

red = blue if and only if *being red = being blue*

but it is hardly illuminating. Rather than explaining the concepts, the truth conditions presuppose the concepts. We might just as well define the “identity condition” of a physical object to be the condition of *being that object* and then claim that physical objects are individuated by their identity conditions. That is about as unilluminating as a theory can be. Unlike the spatio-temporal continuity theory of object individuation, it does not make a substantive claim.

Typically, philosophical logicians slide back and forth between the vacuous claim that concepts are individuated by their truth conditions and the considerably more contentious claim that concepts can be informatively characterized by (and only by) giving truth condition analyses of them. A truth condition analysis of a concept is a *definition* of the concept—an informative statement of necessary and sufficient conditions for something to exemplify the concept. We think it is fair to say that many philosophical logicians do not clearly distinguish between the vacuous claim and the contentious claim, or at least take the vacuous claim to somehow directly support the contentious claim. But we see no reason to think there is any connection between the two claims.

The simplest objection to the truth-condition-analysis theory is that

most concepts do not have the kind of definitions required by the logical theory of concepts. Analytic philosophy in the mid-twentieth century concerned itself almost exclusively with the search for such definitions, and if we can learn anything from that period it is that the search was largely in vain. It is a very rare concept that can be given an informative definition stating truth conditions. This may seem surprising in light of the fact that dictionaries purport to give definitions, and all of the concepts investigated by analytic philosophers have dictionary entries. However it is illuminating to actually consider such a dictionary definition. One dictionary we consulted defined “horse” as “a large four-legged animal, domesticated for carrying riders and hauling loads.” Whatever this definition is, it is not a statement of logically necessary and sufficient conditions for being a horse. For example, a horse does not cease being a horse if it loses a leg in an accident, or if it has never been ridden and used for hauling loads. And if the conditions enumerated by this definition were sufficient for being a horse, then camels would be horses as well. One might suppose that the lexicographers who wrote this definition just did a poor job, but we defy the reader to find a better definition. The only conditions that seem logically necessary for being a horse are very general ones like “occupies space” and perhaps “living creature”, but these are far from adequate to distinguish horses from other animals. The real lesson to be learned from this is that dictionary definitions are not statements of logically necessary and sufficient conditions. Whatever they are, they do not provide the kind of analyses required by the truth condition theory. The importance of this simple objection cannot be overemphasized. Most concepts do not have definitions in the philosophical sense of logically necessary and sufficient conditions. For reasons we find mysterious, many philosophers seem to just ignore this and go on pretending that some form of the truth condition theory of concepts is correct.

4.4.3 The Logical Theory of Concepts

There is another strand to this story. Traditionally, the only logical relations between concepts that were recognized by philosophers were entailment relations. Concepts, as “logical items”, were supposed to be individuated by their logical properties, and it seemed that the only logical properties concepts possessed were those definable in terms of their entailment relations to other concepts. This generates the picture of a “logical space” of concepts, the identity of a concept being determined by its position in the space, and the latter being determined by its entailment relations to other concepts. The claim that concepts must have definitions is just a more specific version of this general picture—one alleging that the position of a concept in logical space is determined not just by one-way entailments but by two-way logical equivalences. Some version of this picture has been prevalent throughout much of twentieth century philosophy, and it still plays a prominent role in philosophical logic. We will call this general picture of the individuation of concepts

the logical theory of concepts. It has often been either confused with or identified with the truth condition theory.

The logical theory of concepts is subject to a rather deep epistemological problem. In general, the logical theory cannot make sense of reasons. To see this, let us begin with defeasible reasons. The logical theory appears to lead directly to the impossibility of defeasible reasons. We assume that what makes something a good reason for holding a belief is a function of the content of the belief. If the content of the belief is determined by entailment relations, then those entailment relations must also determine what are good reasons for holding that belief. The only kinds of reasons that can be derived from entailment relations are reasons that are themselves entailments—conclusive reasons. Thus we are forced to the conclusion that all reasons must be entailments. But this must be wrong, because we have seen that many epistemological problems cannot be solved in terms of conclusive reasons. Justified belief makes essential appeal to defeasible reasoning.

We might try distinguishing between “formal reasons” that derive from principles of logic and apply equally to all concepts, and “substantive reasons” that are specific to individual concepts and reflect the contents of those concepts. The preceding argument is really only an argument that the logical theory of concepts is incompatible with there being non-conclusive substantive reasons. Thus we could render the logical theory of concepts compatible with defeasible reasoning if it could be maintained that all legitimate defeasible reasons are formal reasons. The only plausible way of defending this claim is to maintain that the only legitimate defeasible reasons are inductive reasons and to insist that inductive reasons are formal reasons. This is to take induction to be a species of logic. On this view, there are two kinds of logic—deductive and inductive—and each generates formal reasons that pertain to all concepts and hence need not be derivable from the contents of individual concepts. For example, a conjunction ($P \ \& \ Q$) gives us a reason for believing its first conjunct P regardless of what P and Q are. Similarly, it was traditionally supposed that inductive reasons are formal reasons pertaining equally to all concepts. This absolves us from having to derive inductive defeasible reasons from the essential properties of the concepts to which the reasons apply.

Unfortunately, this attempt to render the logical theory of concepts compatible with induction fails. It was pointed out in chapter one that induction does not apply equally to all concepts. Inductive reasoning must be restricted to projectible concepts. There is no generally accepted theory of projectibility, but it is generally recognized that what makes a concept projectible is not in any sense a “formal” feature of it. The simplest argument for this was given long ago by Nelson Goodman (1955). Define:

x is *grue* if and only if either (1) x is green and first examined before the year 2000, or (2) x is blue and not first examined before the year 2000.

x is *bleen* if and only if either (1) x is blue and first examined before the year 2000, or (2) x is green and not first examined before the year 2000.

'Grue' and 'bleen' are not projectible. For example, if we now (prior to the year 2000) examine lots of emeralds and find that they are all green, that gives us an inductive reason for thinking that all emeralds are green. Our sample of green emeralds is also a sample of grue emeralds, so if 'grue' were projectible then our observations would also give us a reason for thinking that all emeralds are grue. These two conclusions together would entail the absurd consequence that there will be no emeralds first examined after the year 2000. It follows that 'grue' is not projectible. Now the thing to notice is that 'blue' and 'green' are definable in terms of 'grue' and 'bleen' in the precisely the same way 'grue' and 'bleen' were defined in terms of 'blue' and 'green':

x is green if and only if either (1) x is grue and first examined before the year 2000, or (2) x is bleen and not first examined before the year 2000.

x is blue if and only if either (1) x is bleen and first examined before the year 2000, or (2) x is grue and not first examined before the year 2000.

Thus the *formal* relationships between the pair 'blue', 'green' and the pair 'grue', 'bleen' are symmetrical, and hence we cannot distinguish the projectible from the nonprojectible by appealing only to formal properties of the concepts. Projectibility seems to have essentially to do with the content of the concepts. Therefore, any explanation for the existence of inductive defeasible reasons must make reference to the particular concepts to which the reasons apply, and hence, on the logical theory of concepts, inductive defeasible reasons become as mysterious as any other defeasible reasons.

There is of course the further point, defended earlier, that epistemology requires more defeasible reasons than just inductive ones. Thus even if inductive reasons had turned out to be formal reasons, that would not entirely solve the problem of the possibility of defeasible reasons.

The next thing to notice is that the logical theory of concepts makes conclusive reasons just as mysterious as defeasible reasons. This has generally been overlooked, but it is really rather obvious. Epistemologists have noted repeatedly that logical entailments do not always constitute reasons. Some entailments are conclusive reasons and others are not reasons at all. The latter is because P may entail Q without the connection between P and Q being at all obvious. For example, mathematicians

have proven that the Axiom of Choice entails Zorn's Lemma. These are abstruse mathematical principles apparently dealing with quite different subject matters, and just looking at them one would not expect there to be any connection between them. If, without knowing about the entailment, one were so perverse as to believe Zorn's lemma on the basis of the Axiom of Choice, one would not be justified in this belief. Once the entailment is known, you can become justified in believing Zorn's Lemma *partly* by appeal to the Axiom of Choice, but your full reason for believing Zorn's Lemma will be the conjunction of the Axiom of Choice and the proposition that if the Axiom of Choice is true then Zorn's Lemma is true. You are believing Zorn's Lemma on the basis of this conjunction rather than just on the basis of the Axiom of Choice. You can never become justified in believing Zorn's Lemma on the basis of the Axiom of Choice alone, so the latter is not a reason for the former.

On the other hand, some entailments do provide reasons. If I justifiably believe both P and $(P \rightarrow Q)$, I *can* justifiably believe Q on the basis of these other two beliefs. In this case I do not have to believe Q on the basis of the more complicated belief:

P and $(P \rightarrow Q)$ and if $[P \ \& \ (P \rightarrow Q)]$ then Q .

To suppose that each instance of reasoning in accordance with *modus ponens* must be reconstructed in this way would lead to an infinite regress.¹²⁰ Thus some entailments are conclusive reasons and others are not. But the logical theory of concepts gives us no way to make this distinction. It characterizes concepts in terms of their entailment relations to other concepts, but, *a fortiori*, all entailment relations are entailment relations. There is nothing about the entailment relations themselves that could make some of them reasons and others not. Thus conclusive reasons become just as mysterious as defeasible reasons on the logical theory of concepts. This seems to indicate pretty conclusively that the logical theory of concepts is wrong. There has to be more to concepts than entailment relations.

4.4.4 Rational Roles

To argue that the logical theory of concepts is wrong is not yet to say what is right. The theory we want to endorse in its place is the epistemological theory of concepts. This theory begins by noting that concepts are both logical and epistemological items. That is, concepts are the categories whose interrelationships are studied by logic, and they are also the categories in terms of which we think of the world. The interrelationships studied by logic can all be reduced to entailment relations. Thus logic need not take note of any other features of concepts. Logic can get along with a cruder picture of concepts than can epistemology.

120. This was apparently first noted by Lewis Carroll (1895).

But a complete account of concepts must accommodate both logic and epistemology. There is good reason to think that the role of concepts in epistemology is fundamental. Not all entailment relations are conclusive reasons, but it seems likely that all entailment relations derive from “simple” entailment relations, where the latter are just those that are conclusive reasons. Thus a theory of concepts adequate for epistemology will very likely be adequate for logic as well. The question then becomes, “What kind of theory of concepts is adequate for epistemology?”

In epistemology, the essential role of concepts is their role in reasoning. Concepts are the categories in terms of which we think of the world, and we think of the world by reasoning about it. This suggests that concepts are individuated by their role in reasoning. What makes a concept the concept that it is is the way we can use it in reasoning, and that is described by saying how it enters into various kinds of reasons, both conclusive and *prima facie*. Let us take the *rational role* of a concept to consist of (1) the reason-schemas (conclusive or defeasible) licensing an inference to the conclusion that something exemplifies it or exemplifies its negation, and (2) the reason-schemas licensing conclusions that can be justifiably drawn (conclusively or defeasibly) from the fact that something exemplifies the concept or exemplifies the negation of the concept.¹²¹ We have encountered reason-schemas throughout this book, and we are able to use them to individuate perceptual concepts such as red. For instance,

“S appears red to me” is a defeasible reason for me to believe that S is red.

In our view, the concept red is individuated in part by the fact that a certain kind of belief is licensed by the defeasible reason-schema that applies to it (we will explore various other reason-schemas in chapter seven). Taken together, the reason-schemas we use to think about the world constitute our epistemic norms. Thus, the epistemic norms governing a concept are descriptive of its rational role.

Our proposal is that concepts are individuated by their rational roles. The essence of a concept is to have the rational role that it does. If this is right, the explanation for how there can be such things as defeasible reasons becomes trivial. Defeasible reasons are primitive constituents of the rational roles that characterize concepts. Defeasible reasons need not have an origin in something deeper about concepts, because there is nothing deeper. In an important sense, there is nothing to concepts over and above their rational role. To describe the rational role of a concept is to give an analysis of that concept, although not a truth condition analysis.¹²²

121. We are leaving out some subtleties, as they are not particularly relevant for the problem of relativism that we are presently trying to solve. For a lengthier discussion of rational roles, see Pollock (1989), chapters four and five.

122. This view of concepts is reminiscent of the verification theories of the logical

It should be noted that the rational role account of concept individuation is only distantly related to proposals in the philosophy of mind that fall under the label of *conceptual role semantics*. These are theories that claim that the nature of the meaning of a thought consists in the (typically inferential) relations that the thought has to other thoughts.¹²³ It has been frequently asserted that this view is committed to meaning holism, in the sense that that two people who have different beliefs and draw distinct inferences based on those beliefs will have different thoughts.¹²⁴ Our rational role semantics focuses only on the reason-schemas that guide reasoning. The beliefs that provide the premises for the reasoning are irrelevant to rational roles.

We think it is undeniable that concepts are individuated by their rational roles, and not (at least in any non-vacuous way) by their truth conditions. But some further explanation for all of this is required. *Why* are concepts individuated in this way? We will shortly propose an answer to this question. For the moment, however, we will simply take it as established that concepts are individuated in this way. The importance of this theory of concepts for the matters at hand is that it lays to rest the spectre of epistemological relativism. Epistemological relativism is the view that (1) different people could have different epistemic norms that conflict in the sense that they lead to different assessments of the justifiedness of the same belief being held on the same basis, and (2) there is no way to choose between these norms. The epistemological theory of concepts enables us to escape any such relativism. Because concepts are individuated by their rational roles, it becomes impossible for people's epistemic norms to differ in a way that makes them conflict with one another. The epistemic norms a person employs in reasoning determine what concepts she is employing because they describe the rational roles of her concepts. If two people reason in accordance with different sets of

positivists. Pollock first defended a theory of this sort in his (1968), and in more detail in his (1974), although in those publications he talked about "justification conditions" rather than rational roles, and used the term a bit more narrowly. This view of concepts is also related to the somewhat cruder views expressed by Michael Dummet (1975) and (1976) and Hilary Putnam (1979) and (1984).

123. See Hartry Field (1977) and Ned Block (1986). Under pressure from Hilary Putnam's Twin Earth cases (1975), many conceptual role theorists have defended so-called 'two factor' versions of the view. An internal factor fixes the narrow content of a thought for use in psychological explanation, and an external factor—typically a theory of truth—fixes the wide content of a thought for resolving questions of reference.

124. The most vocal critics who pursue this line are Jerry Fodor and Ernest LePore (1992). See also Rob Cummins (1989) and (1996). Cummins comes closest to offering an objection that might be applied to our rational role account when he alleges that reason-schemas are to be revealed by psychological investigation and that psychology requires that concepts be fixed first (1996, 43). Our account of reason-schemas will evade this criticism if we offer a methodology for determining our reason-schemas that does not require psychological explanations. In the next chapter, we offer just such a methodology.

epistemic norms, all that follows is that they are employing different concepts. Thus it is impossible for two people to employ different epistemic norms in connection with the same concepts. Their conceptual frameworks are determined by their epistemic norms. Epistemological relativism is logically false.

We have argued that if two people use different epistemic norms, then they are employing different concepts and this foils relativism. It might be wondered, finally, whether or not people actually use different epistemic norms. We doubt that there really is any variation in epistemic norms from person to person.¹²⁵ We suspect that epistemic norms are species-specific, but this is an empirical question.¹²⁶ In order to resolve these issues, what seems to be required is a methodology for determining what our epistemic norms are. We turn to that problem in the next chapter.

5. Conclusions

The main purpose of this chapter has been to understand how epistemic norms function. They guide us in our cognition, and beliefs are justified just in case they are held in compliance with epistemic norms, but the way in which they guide us proves more difficult to understand than epistemologists have often supposed. Many epistemologists have been tempted by the intellectualist model, according to which we make explicit appeal to epistemic norms. But the intellectualist model could not be a correct theory of the way epistemic norms function, because it would lead to an infinite regress. In order to comply with an epistemic norm, we would have to have a justified belief to the effect that the norm makes a certain prescription in the present case, and that would require us to comply with another epistemic norm. The principal insight of the chapter is that epistemic norms function in the same general way as other procedural norms. They are descriptive of our procedural knowledge of how to cognize, articulating what we “know to do” in cognizing. As with any procedural norms, we do not always succeed in complying with them, so there is a competence/performance distinction in epistemology just as there is in linguistics. A theory of the content of our epistemic norms is not just a description of what we do when we cognize.

125. This possibility is one of the central concerns in Stephen Stich (1992).

126. The conclusion that if different people employ different epistemic norms then they employ different concepts may seem puzzling because it appears to make it inexplicable how such people could communicate with each other. Even if our conjecture regarding the species-specificity of norms is false, it need create no difficulty for communication. Pollock has argued at length that concepts play only an indirect role in communication. (Pollock’s entire theory of language is developed in his (1982). A briefer sketch of the theory can be found in chapter two of Pollock (1984). The reader who is concerned with this question should consult those books.)

That would be a performance theory. It is instead a competence theory, describing the way we know how to cognize, whether we actually do it that way or not.

Because epistemic norms are internalized, they must be able to function without conscious monitoring. This has two important consequences. First, epistemic norms must appeal only to internal states. Second, epistemic norms can appeal to any internal states, not just beliefs. So we simultaneously have a refutation of belief externalism and an explanation for why the doxastic assumption fails.

Norm externalism alleges that epistemic norms can be evaluated in terms of external properties like their reliability. It turns out, however, that such evaluations are already built into our actual epistemic norms. As such, external considerations cannot mandate changes in our norms. Apparently our epistemic norms are beyond criticism. This seems initially puzzling, but it is explained by endorsing a rational role theory of concept individuation.

The preceding remarks explain how epistemic norms work, but they do not determine which epistemic norms are correct. We turn next, then, to the methodology for determining the correct norms and finally, in chapter seven, to the norms themselves.