IV

The Structure of Epistemic Justification

JOHN L. POLLOCK

THE concept of epistemic justification, or of a justified belief, is one of the most important concepts of philosophy. The fundamental problem of epistemology is one of answering questions of the form "How do you know that *P*?," and these latter questions have customarily been interpreted as questions concerning what justifies one in believing that *P*. But it is my contention that philosophers have generally worked with a grossly over-simplified picture of the structure of epistemic justification. They have commonly accepted a deductive model in which a reason can only be a reason by virtue of logically entailing what it is a reason for. This deductive model has been largely responsible for generating a number of traditional philosophical problems and making them appear insoluble. For example, there is no deductive connection between perception and the material world, or behavior and mental states, or inductive grounds and the inductive conclusion. Consequently, accepting the deductive model it is impossible to explain how we can have knowledge of the material world, other minds, or contingent generalizations. In this paper the attempt will be made to clarify as much as possible the way in which justification actually works, showing that it is much more complicated than the deductive model would have us believe, and showing that once the true structure of epistemic justification is understood, many traditional problems appear relatively unproblematic.

I. Reasons

What is it that justifies a belief? Suppose someone justifiably believes some fact about the world on the basis of some other fact. Philosophers have often wanted to say that it is the second fact that justifies one's belief in the first fact. For example, consider the case of a person who believes there is a sheep in the field because he sees a dog that looks very much like a sheep—so much like a sheep in fact that anyone would be justified in taking it to be a sheep until he

examined it quite closely. One is apt to say that it is the *fact* that the dog looks like a sheep that justifies the person in thinking that there is a sheep in the field. But this is misleading. What is important in deciding whether the person is justified in his belief is not the fact itself, but rather the person's belief that it is a fact. After all, if the person did not believe that the dog looked like a sheep, then his belief that there was a sheep in the field would not be justified, although it would of course still be a fact that the dog looked like a sheep. Thus we must say that what justifies a belief is always another belief. It is a person's "doxastic state" which determines which of his beliefs are justified. Of course, we can still talk about facts, states of affairs, etc., justifying beliefs, but this must be understood in terms of beliefs justifying beliefs.

In order to justify a belief one must appeal to another belief, but clearly the simple fact that a person believes one thing does not automatically mean that he is justified in believing something else which can be supported by the first belief. He must not only *have* the first belief—he must also be *justified* in having it. If a person believed, for no good reason, that the moon was shaped like half of an egg shell, so that it had no back side, this belief would not justify his believing further that the pictures of the back side of the moon obtained by lunar probes were fraudulent. Thus to justify a belief, one must appeal to a further justified belief.

Justification proceeds in terms of reasons. When one belief justifies another, then the former is said to be a *reason* for the latter.[1] Sometimes, rather than talk about *S*'s belief-that-*P* being a reason for *S* to believe that *Q*, we say that *S*'s reason is the fact that *P*. But, for the reasons just given, talk about beliefs as reasons seems to be primary.

Reasons differ in what we might call "quality." One belief may be a better reason for believing something than another belief. For example, if one has tested a hypothesis in ten different cases this might be sufficient to justify him in believing the hypothesis true, and so would provide him with a reason; but if he had tested the hypothesis in a hundred different cases, this would surely provide him with a better reason for believing the hypothesis true. It seems that all reasons are susceptible to these differences in quality. It is clear, for example, that one inductive reason may be better than another

[1] This is not intended to be a definition of "reason," but only a rather loose characterization. As will become apparent in this paper, the concept of a reason is really a family of concepts which share many family resemblances but about which little can be said in general. No attempt will be made to give a general definition of "reason," although definitions will be offered for certain important classes of reasons.

because it is based on a larger sample. But even in the case of perceptual beliefs there are differences. For example, if an object seen in broad daylight looks plainly red to me, this may justify me in being certain that it is red. But if the object is seen in a dark room where I can just barely discern colors, and it looks only vaguely red to me, this may only justify me in thinking that it is probably red.

Let us say that a *good* reason is one that is sufficient to justify the belief for which it is a reason. We often have reasons both for believing something and for disbelieving it. These are all reasons, both pro and con, but they do not all justify what they are reasons for. Each one by itself, in the absence of any competing reasons to the contrary, would be a good reason for believing what it is a reason for, but taken together with the other reasons, it may no longer be a good reason. Thus reasons need not always be good reasons.

II. The Perspicuity of Reasons

In order for the belief-that-*P* to be a good reason for a person *S* to believe that *Q*, he must know that it is. Reasons must be *perspicuous*. Why is this so? First, suppose that *S* believed that *P* and believed that *Q*, but did not think that *P* was a good reason for believing that *Q*. Then clearly for him to believe that *Q on the basis of P*, while simultaneously denying that he believes *P* to be a good reason for believing that *Q*, would be irrational. He would not be justified in believing that *Q*, and so *P* would not be a good reason for him to believe that *Q* (because good reasons justify). Thus it seems that:

(I) If *P* is a good reason for *S* to believe that *Q*, then it must be the case that *S* believes that it is.

In apparent contradiction to (I), we sometimes talk about a person having a good reason for believing something without realizing it. For example, we may say of a detective that the clues he has found constitute a good reason for thinking the gardener did it, but that he does not realize this (he has not drawn the proper conclusions). But in a case like this, the clues are not really a good reason for *him* (the detective) to believe that the gardener did it, because he has not seen the connection. It would be unreasonable for him to believe on the basis of the clues that the gardener did it unless he had seen the connection. What *is* true is that the clues *would be* a good reason for him to believe that the gardener did it *if* he saw the connection. Talk about unrecognized good reasons is elliptical for talk about what would be good reasons if people saw the right connections.

An objection that is apt to arise to principle (I) concerns those beliefs that are based directly upon perception. For example, I would urge that frequently a person's reason for believing that something is red is the fact that it looks red to him. But does he believe that this is his reason? This might seem suspect, because this is not an answer that a non-philosopher would be apt to give to the question "Why do you believe that that is red?" But in fact, I think that a person does believe that that is his reason. At least, if we asked him "Do you believe that it is red because it looks red to you (as opposed, e.g., to because you were told that it is red)?," he would no doubt answer in the affirmative. On the other hand, if you simply asked him "Why do you think it is red?," without suggesting the answer "Because it looks red to me," he might not know what to say. But I do not think that this shows he does not believe that he thinks it is red because it looks red to him. It is just that he does not know what you are getting at with your question, because the answer is so *obvious*.

Next, suppose that *S* believes (1) that *P*, believes (2) that *Q*, and believes (3) that *P* is a good reason for him to believe that *Q*, but is not justified in believing (3). For example, a person might believe that extrasensory perception is a good reason for thinking that the next card he turns up will be an ace, but not be justified in that belief. Under these circumstances it seems clear that he would not be justified in that belief. Under these circumstances it seems clear that he would not be justified in thinking that the next card will be an ace, and thus his extrasensory perception would not constitute a good reason. Therefore:

(II) If the belief-that-*P* is a good reason for *S* to believe that *Q*, then *S* must justifiably believe that it is.

But not only is this the case, but also the converse, viz., if *S* believes, and is justified in believing, that the belief-that-*P* is a good reason for him to believe that *Q*, then it is a good reason. For suppose that a person does justifiably believe that the belief-that-*P* is a good reason for him to believe that *Q* (e.g., that extrasensory perception is a good reason for believing that the next card will be an ace). Then clearly if he justifiably believes that *P*, he is justified in believing that *Q*, and thus *P* is a good reason (because it justifies). So:

(III) If *S* justifiably believes that the belief-that-*P* is a good reason for him to believe that *Q*, then it is a good reason for him to believe that *Q*.

Therefore *S*'s justified belief that the belief-that-*P* is a good reason for him to believe that *Q makes it true* that *P* is a good reason, and so cannot be wrong. Such a belief is not *mere* belief, but knowledge. Thus:

(IV) If the belief-that-*P* is a good reason for *S* to believe that *Q*, then *S* must know that it is.

Reasons are perspicuous.

III. Logical and *A Posteriori* Reasons

There are two fundamentally different kinds of good reasons for believing things. Some beliefs are good reasons for believing other things simply by virtue of their logical nature. For example, the justified belief that *that* is a sheep in the field is a good reason for one to believe that *there is* a sheep in the field, and this is so simply because of the meanings of the statement that that is a sheep in the field and the statement that there is a sheep in the field (this is not *quite* right, as will be seen shortly). Whenever the justified belief-that-*P* is a good reason for one to believe that *Q*, simply by virtue of the meanings of the statements that *P* and that *Q* we might say (provisionally) that the statement-that-*P* is a *logical reason* for believing the statement-that-*Q*. The simplest examples of logical reasons are of course simple entailments, but it will be seen later that there are other logical reasons that are not entailments.

The relation of one statement to another wherein justified belief in the first provides a logical reason for believing the second is intended to be a logical relation between statements. It results simply from the logical nature of the statements. However, this is somewhat difficult to reconcile with the fact that good reasons must be perspicuous. It seems that whether *P* is a logical reason for *S* to believe that *Q* must depend in part on whether *S* sees that it is. In order to avoid this, we must modify the notion of a logical reason. A necessary condition for a person to understand the statement-that-*Q* is that he knows what would justify him in believing that *Q*. Thus we can say that if *P* is a logical reason for believing that *Q*, but *S* does not see that it is, then *S* simply does not understand the statement-that-*Q*. This must lead us to refine the definition of "logical reason" as follows:

> The statement-that-$P(S)$ is a logical reason for *S* to believe the statement-that-$Q(S)$ if, and only if, necessarily, for any person S' who understands both statements, it is possible for the belief-that-$P(S')$ to be a good reason for S' to believe that $Q(S')$ without his having an independent reason for believing that $P(S') \supset Q(S')$.

It should be noted that in this definition, by writing "$P(S)$" rather than just '*P*' I have allowed for the possibility that some logical reasons may be "subject-relative." This, of course, is not true for entailments, but I will urge shortly that there are other logical reasons for which it is true. For example, it will be argued that "*x* looks red to *S*" is a logical reason for *S* (but not for someone else) to believe that *x* is red.

It should also be noted that the above definition allows for the defeasibility of some logical reasons. That is, the definition does not require that the belief-that-$P(S')$ *always* be a good reason for S' to believe that $Q(S')$. It merely requires that it is *possible* for $P(S')$ to be a good reason for S' to believe that $Q(S')$ without S' having an independent reason for believing that $P(S') \supset Q(S')$, and leaves open the possibility that under *some* circumstances $P(S')$ may not be a good reason for S' to believe that $Q(S')$. This possibility will be discussed in Section V. Again, this possibility does not arise for logical reasons that are entailments.

One thing that should be noted about the definition of "logical reason" is that *P* may be a logical reason for *S* to believe that *Q* without *P* being a good reason for *S* to believe that *Q*. This sounds paradoxical, but it is not. The distinction is roughly that between "a good reason for believing that *Q*" and "a good reason for *S* to believe that *Q*." For example, that two times two equals four is a good reason for thinking that two is the square root of four, but it may not be a good reason for Jones to believe that two is the square root of four because Jones may not know what a square root is. A logical reason for *S* to believe that *Q* is only a good reason for *S* to believe that *Q* if *S* understands the statement-that-*Q*. Let us say that *S has* the logical reason *P* for believing that *Q* if (1) *P* is a logical reason for *S* to believe that *Q*, (2) *S* justifiably believes that *P*, and (3) *P* is a good reason for *S* to believe that *Q*.

Philosophers have always recognized the category of logical reasons. But they have often overlooked the fact that there are other good reasons that are not logical reasons. For example, the belief that Smith said there is a woman in the next room may be a good reason for me to think there is a woman in the next room. But it is certainly not a logical reason. Let us say that a good reason that is not a logical reason is an *a posteriori reason.*

Although there are these two categories of good reasons—logical reasons and *a posteriori* reasons—they are closely related. At least often a belief is only an *a posteriori* reason for a person to believe something if the person also has a related logical reason for believing it. For example, the belief that Smith said there is a woman in the next room is only a good reason for me to think that there is a woman in the next room because I justifiably believe that Smith is telling the truth. And the conjunctive belief that Smith said there is a woman in the next room and he is telling the truth, is a logical reason for me to think that there is a woman there. Thus in many cases, the belief *P* is only an *a posteriori* reason for *S* to believe that *Q* if *S* also justifiably believes some further statement *R* such that the statement-that-(*P* & *R*) is a logical reason for believing the statement-that-*Q* and *S* understands the statement-that-*Q*. Such an *a posteriori* reason only becomes an *a posteriori* reason when one acquires such a further justified belief.

The question arises whether all *a posteriori* reasons can be reduced to logical reasons in this way. It seems that they can. Because reasons are perspicuous, if the belief *P* is a good reason for *S* to believe that *Q*, then he must know that it is. This in turn seems to imply that *S* must justifiably believe that if *P*, then *Q*. For example, if *S* knows that *x*'s looking red to him is a good reason for him to believe that *x* is red, then *S* must justifiably believe that if *x* looks red to him it is red. Analogously, if *S* knows that Smith's having told him that there is a woman in the next room is a good reason for him to believe that there is a woman in the next room, then *S* must justifiably believe that if Smith told him there is a woman in the next room, then there is a woman in the next room. (Of course, *S* need not justifiably believe that *whenever* something looks red to him it is red, or that *whenever* Smith tells him there is a woman in the next room there is a woman in the next room.) The conjunctive belief that (*P*, and if *P* then *Q*) is a logical reason for *S* to believe that *Q*. Consequently, although there are these two distinct classes of good reasons—logical and *a posteriori* —*a posteriori* reasons can always be reduced to logical reasons. Thus justification can be thought of as proceeding exclusively in terms of logical reasons.

It should be observed that to say that *P* is a logical reason for *S* to believe that *Q* merely expresses a logical connection between *P* and *Q* and does not imply that *S* believes that *P*, whereas to say that *P* is an *a posteriori* reason for *S* to believe that *Q* only expresses a contingent connection between *P* and *Q* and does imply that *S* believes that *P*. We might say that the concept of a logical reason is purely relational while that of an *a posteriori* reason is not. The reason for the variance in terminology is that *any* belief *P* could be an *a posteriori* reason for *S* to believe that *Q*, provided only that *S* has suitable grounds for believing that if *P*, then *Q*.

IV. Conclusive and Non-conclusive Logical Reasons

Simple entailments are logical reasons. It is tempting to suppose that *all* entailments are logical reasons. But this will not do, because reasons must be perspicuous. With a fairly complex entailment, one might be justified in believing the antecedent, but—not knowing that the antecedent does entail the consequent—not be justified in believing the consequent. For example, the Axiom of Choice entails Zorn's Lemma, but this is not an entailment that one could be expected to see without proof. Thus one might justifiably believe that the Axiom of Choice is true, but not see that it entails Zorn's Lemma, and so not be justified in believing that Zorn's Lemma is true. In order for an entailment to provide a good reason, one must be justified in believing that there is such an entailment. In the case of simple entailments, we can see just by considering the meanings of the statements that the one entails the other, and so such entailments are always good reasons for anyone who understands the statements involved. But there are more complex entailments that can only be known as the result of giving a demonstration, and these entailments do not provide us with good reasons until we have given the requisite demonstrations. Hence, although after the demonstrations have been given they may become good reasons, they are not logical reasons. They are only *a posteriori* reasons.

We have thus isolated one class of logical reasons, which we can call "conclusive" reasons. Let us say that:

> The justified belief-that-*P* is a *conclusive reason* for *S* to believe that *Q* if, and only if, the statement-that-*P* entails and is a logical reason for *S* to believe that *Q*.

Conclusive reasons are logical reasons. Do conclusive reasons exhaust the class of logical reasons? It is astonishing how often philosophers have thought that they do. For example, Judith Jarvis Thomson, in her article "Reasons and Reasoning," writes: "One thing that seems plain is that if 'This is *S*' does not imply 'This is *P*' it will always be at best a matter of fact, to be established by investigation, that the first is a reason for the second."[2] But there seem to be clear counterexamples to the supposition that all logical reasons are conclusive. Consider, for example, induction. A class of singular statements $P_1, \ldots, P_n$ can provide inductive grounds for an un-

[2] *Philosophy in America*, ed. by Max Black (Ithaca, 1965), p. 292. See also Bertrand Russell, *Human Knowledge: Its Scope and Limits* (London, 1948).

restricted generalization Q, and thus the conjunction (P_1 & . . . & P_n) of those singular statements is a good reason for the general statement Q. The conjunction of the singular statements clearly does not entail the generalization, and yet it seems to constitute a logical reason for believing it. This can be shown as follows. If (P_1 & . . . & P_n) constituted merely an *a posteriori* reason for believing that Q, then a person would have to have an independent reason for believing that $[(P_1 \& \ldots \& P_n) \supset Q]$ before the conjunction could be a reason for believing that Q. There are only two plausible alternatives for what such an independent reason might be. First, it might be an inductive reason. However, this would merely push the difficulty back one step without removing it, so this alternative is no solution. Second, this conditional might be entailed by some general principle which we might call a principle of the Uniformity of Nature. Philosophers have often looked for some such principle to justify induction. But even if such a principle could be found, which seems exceedingly unlikely, we could still ask what justifies us in believing it. Again, there are two possibilities. First, the principle might be a truth of logic. But in that case the conditional $[(P_1 \& \ldots \& P_n) \supset Q]$, being entailed by the general principle, would also be a truth of logic. And this would require that the conjunction of singular statements (P_1 & . . . & P_n) entail the general statement Q, which is impossible. Suppose instead that the general principle is not a truth of logic. Then we can ask what logical reasons there are for believing it. And on the supposition that all logical reasons are conclusive, this amounts to asking what justified beliefs entail it. These beliefs constitute a set Γ. The beliefs in Γ must be justified without appeal to the general principle, so they cannot include any beliefs justified inductively. Consequently, any general beliefs in Γ must be truths of logic. But if a set of statements some of which are truths of logic entails another statement, then the set of statements that results from deleting the truths of logic also entails the other statement. Consequently, any general beliefs in Γ can be omitted with the result that the remaining beliefs still entail the principle of the Uniformity of Nature. The remaining beliefs in Γ must constitute a class of singular statements $R_1, \ldots, R_m$. But then if the inductive grounds $P_1, \ldots, P_n$ together with the principle of the Uniformity of Nature entailed Q, the inductive grounds $P_1, \ldots, P_n$ together with the additional singular statements $R_n, \ldots, R_m$ would also have to entail Q (because the latter entail the principle of the Uniformity of Nature). And then again we would have a general statement being entailed by a finite conjunction of singular statements, which is impossible. It seems then

that induction provides us with an example of a non-conclusive logical reason. And it will be argued shortly that there are others.[3]

The supposition that all logical reasons are conclusive ones is related to the traditional hope that all epistemological problems can be solved by giving reductive analysis. Philosophers have been enamoured with the idea that you can order all statements into an infinite sequence such that each statement in the sequence is either epistomologically basic (and hence not in need of justification)[4] or else can be expressed as a logical construction of statements preceding it in the sequence. Let us call such a sequence a "reductive sequence." Such a sequence would provide us with reductive analyses of all statements in terms ultimately of epistemologically basic statements. Thus, for example, statements about material objects were supposed to be analyzed in terms of statements about sense-data (phenomenalism), statements about other minds were supposed to be analyzed in terms of statements about bodies (behaviorism), ethical statements were supposed to be analyzed in terms of nonethical statements (naturalistic ethics), and so on.

On the supposition that such reductive analyses were always possible, it became plausible to suppose that all logical reasons were conclusive. One could think of the chain of justification as proceeding up through successively higher levels of the reductive sequence in such a way that each statement in the chain was a logical construction out of earlier statements in the reductive sequence and thus entailed by them. Then ultimately a justified belief would be entailed by the epistemologically basic statements at the bottom of its chain of justification. But in recent years it has come to seem quite unlikely that this enterprise of seeking reductive analyses can be successful. It does not seem to be the case that reductive analyses can generally be given. And it becomes accordingly less likely that justification can always be viewed in terms of entailment. It appears that there must be logical reasons that are not conclusive ones.

[3] This of course should not be taken as implying that induction and other non-conclusive logical reasons cannot be "reduced" to logical reasons, utilizing the techniques of Section III, by conjoining them with a conditional stating that if the reason is true then the conclusion is true. But the point is that these reasons remain logical reasons even without conjoining them with such a conditional.

[4] The traditional view of epistemology, which I accept, is that there are certain beliefs which do not stand in need of justification, and which in turn constitute the basis upon which all other beliefs are ultimately justified. These beliefs are said to be *epistemologically basic*. Various candidates have at different times been proposed as being epistemologically basic—sense datum beliefs, introspective beliefs, etc.—but for present purposes there is no need to decide what beliefs are epistemologically basic, or what properties epistemologically basic beliefs have. In fact, it is in no way essential to this paper that there are any such beliefs.

V. Logically Good Reasons and *Prima Facie* Reasons

Logical reasons that are not conclusive are particularly interesting because they have been largely overlooked by philosophers bent on finding reductive analyses. Perhaps in them lies the key to a number of stubborn epistemological problems. Let us call such reasons *logically good reasons*. Whereas conclusive reasons guarantee truth, logically good reasons only guarantee justification. Induction provides us with one example of a logically good reason. Clearly the inductive grounds for a conclusion do not constitute a conclusive reason for believing the conclusion, but as was argued above, they do constitute a logical reason.

Many logically good reasons have a certain kind of structure which makes it reasonable to call them *prima facie reasons*. A *prima facie* reason is a reason that by itself would be a good reason for believing something, and would ensure justification, but may cease to be a good reason when taken together with some other beliefs. In other words, a *prima facie* reason is a logically good reason that is defeasible.

In order to explore the structure of *prima facie* reasons, let us seek several examples of them. Induction provides us with one example. As we have seen, an inductive reason is a logically good reason, and it is clearly defeasible. An inductive generalization can be rebutted on at least two grounds. First, no matter how strong the initial inductive evidence for the generalization, if further investigation turns up a counterexample then the original reason ceases to be a good reason. Second, if it is discovered that the sample on which the original generalization was based was not a fair sample, this will make the initial reason no longer a good reason even though it was a good reason until this was discovered.

Another, perhaps more interesting example, concerns perceptual judgments. I will consider only one instance of a perceptual judgment —a person's judgment that something is red on the basis of its looking red to him—but it seems clear that the conclusions drawn in connection with this one example have rather broad application to perceptual judgments in general. It seems indisputable that there must be some sort of logical relation between "x looks red to S" and "x is red." It is not just an accident that red things tend to look red to people. This vague intuition is fortified by the observation that to suppose otherwise would make it impossible for us to ever know that anything is red. If we were to suppose that the connection between something's looking red to us and its actually being red is only a contingent connection, then the only way we could ever establish the

connection is inductively. But we could never establish inductively that things that look red to us tend to be red, because in order to do that we would have to be able to tell independently what things are red, and the only way we have of doing that is in terms of what things look red to us, which would beg the question. Therefore, if the connection were merely contingent, then knowledge of red objects would be impossible. But knowledge of red objects is of course possible, so the connection cannot be contingent; it must be a logical connection of some sort.

The hope that this logical connection could be explained entirely in terms of conclusive reasons was just the hope of phenomenalism. It was hoped that the meaning of "x is red" could be analyzed in such a way that x's looking red under certain specifiable phenomenological conditions would logically entail that x is red. I think it is fair to say that we know now that no such analysis can be given. Consequently this logical connection cannot be explained entirely in terms of conclusive reasons. But given that it is a *logical* connection, it must be explained in terms of some sort of logical reasons, and hence these reasons must be logically good reasons.

Can we describe the logical connection between "xl ooks red" and "x is red" in such a way as to elicit the structure of the logically good reasons involved? Ordinarily, when I can see an object clearly, and have no reason for supposing either that there is something wrong with my eyes, or that there are strange lights playing on the object, or anything of that sort, I unhesitatingly judge that the object is red if it looks red to me. If I do not have any beliefs about x other than that it looks red to me, then I am justified in thinking that it is red, and this is so simply by virtue of the concepts "red" and "looks red." But if I do have certain other beliefs, my belief that x looks red to me may not justify me in believing that x is red. For example, I may believe that there are red lights shining on x, and that in the daylight it looks white. If I had those beliefs, the simple fact that x looked red to me would not justify me in believing that x was red. Thus the belief that x looks red to me is a defeasible logically good reason for me to think that x is red, i.e., it is a *prima facie* reason.

In order to clarify the way in which a *prima facie* reason looks, we must talk about *excluders*—the beliefs which when conjoined with the *prima facie* reason may prevent it from justifying the belief for which it is a *prima facie* reason. There are two kinds of excluders. First, if P is a *prima facie* reason for a person S to believe that Q, then any reason for S to think that $\sim Q$ is an excluder. This is the simplest kind of excluder. Let us call it a *type I* excluder. For example, "Jones told me that x is not red, and Jones is generally reliable" would be an example of a type I excluder for "x looks red to me" as a *prima facie* reason for me to believe that x is red. Analogously, "That crow is not black" would be a type I excluder for an inductive reason for thinking that all crows are black.

There is also a second kind of excluder. For example, although the belief that there are red lights shining on x is not a reason for thinking that x is not red, it is nevertheless sufficient to prevent the *prima facie* reason that x looks red to me from justifying the belief that x is red, and hence it is an excluder. This second kind of excluder is, roughly speaking, a reason for thinking that in this kind of case, knowing-that-P is not a good way to find out whether Q. For example, if there are red lights shining on x, then knowing that x looks red to me is not a good way to find out whether x is red, because the red lights can make a white object look red. More precisely, if P is a *prima facie* reason for Q, then R is an excluder of this second kind just in case it is a reason for one to believe that the subjunctive conditional, "If it were not true that Q, then it would not be true that P," is false (thus, e.g., "There are red lights shining on x" is a reason for thinking that the conditional "If x were not red, then it would not look red to me" is false). "If it were not true that Q, then it would not be true that P" is actually the contrapositive of the conditional I have in mind, but it is difficult to state in any other way without making it sound like either a material conditional or a counterfactual conditional. Henceforth this subjunctive conditional will be written simply as "$P \rightarrow Q$." Thus the second kind of excluder is a reason for thinking that $\sim(P \rightarrow Q)$. Let us call these *type II* excluders.

In the case of induction, a type II excluder is just a reason for thinking that the inductive sample is not a fair sample. This is because what is meant by saying that the sample is not a fair sample is just that examining it does not constitute a good way of finding out whether the predicate in question is universally satisfied. For example, if one is attempting to show inductively that no automobile can attain a maximum speed greater than eighty miles per hour, but it is subsequently discovered that all the automobiles examined were Volkswagens and Volkswagens have less power than many other automobiles, this constitutes a type II excluder for the inductive generalization.

If P is a *prima facie* reason for S to believe that Q, then in order for S's justified belief-that-P to justify him in believing that Q, he must have no reasons for believing either that $\sim Q$ or that $\sim(P \rightarrow Q)$. But in fact, S's simply not *having* any such excluding reasons is not enough. It must also be the case that he *does not believe* that he has any such excluding reasons. For example, suppose that S is trying to predict the colors of marbles drawn from an urn. Let us suppose that fifteen marbles have been drawn so far, and they have all been red. He might then conclude inductively that the next marble will also be red. If he has no reasons for thinking either that the next marble will not be red, or that there is something fishy about the urn so that the inductive generalization does not provide him with a way of getting to know whether the next marble will be red, then it seems we would want to say that he would be justified in believing that the next marble will be red. In other words, he has a *prima facie* reason for thinking that the next marble will be red. But in fact, this is not enough. Let us suppose that S believed, without justification, that he had ESP and that on that basis he could tell that the next marble would be black. Even though his belief in this excluder is not justified, the mere fact that he does believe it would be enough to prevent him from being justified in believing that the next marble will be red.

Therefore, if a *prima facie* reason justified S in believing that Q, then it must be the case both that S *has* no excluding reasons and that he *does not believe* that he has any excluding reasons. But this can be simplified somewhat, because if S does not believe he has any excluding reasons, then it follows that he does not have any. This is because, as has already been argued, in order for something to be a reason for S he must believe that it is. So we can say simply that:

> If P is a *prima facie* reason for S to believe that Q, and S justifiably believes that P and does not believe that he has any reasons for thinking either that $\sim Q$ or that $\sim(P \rightarrow Q)$, then S is justified in believing that Q.

The above principle can be reformulated to yield a definition of "*prima facie* reason." This cannot be done directly without circularity because "reason" is used in the statement of the principle, and one kind of reason that might be involved is another *prima facie* reason. However, this circularity can be easily avoided by recalling that any reason can be replaced by a conclusive reason, and hence we need only worry about whether S has any conclusive reasons for thinking either that $\sim Q$ or $\sim(P \rightarrow Q)$. Given this, a noncircular definition of "*prima facie* reason" can be formulated:

> The statement-that-$P(S)$ is a *prima facie reason* for S to believe that $Q(S)$ just in case $P(S)$ is a logical reason for S to believe that $Q(S)$ and it is necessarily true that, for any person S' who understands both statements, (1) is S' justifiably believes that $P(S')$ and does not

believe that he has any conclusive reasons for thinking either that $\sim Q(S')$ or that $\sim[P(S')\rightarrow Q(S')]$, then S' is justified in believing that $Q(S')$, and (2) it is possible for S' to have a conclusive reason for thinking that $\sim Q(S')$ or that $\sim[P(S')\rightarrow Q(S')]$.

Actually, this definition can be simplified a bit. The two kinds of excluders can be collapsed into a single kind. This is because, given that S justifiably believes that P, a conclusive reason for him to believe that $\sim Q$ is also an *a posteriori* reason for him to believe that $\sim(P\rightarrow Q)$. This is because the conjunction $(P \,\&\, \sim Q)$ is a conclusive reason for believing that $\sim(P\rightarrow Q)$. Thus (assuming that S understands these statements) whenever S thinks he has a reason for believing both that P and that $\sim Q$, he must also think he has a reason for believing that $\sim(P\rightarrow Q)$. Consequently the above characterization of *prima facie* reasons can be simplified to read:

> The statement-that-$P(S)$ is a *prima facie* reason for S to believe that $Q(S)$ just in case $P(S)$ is a logical reason for S to believe that $Q(S)$ and it is necessarily true that, for any person S' who understands both statements, (1) is S' justifiably believes that $P(S')$ and does not believe that he has any conclusive reasons for thinking that $\sim[P(S')\rightarrow Q(S')]$, then S' is justified in believing that $Q(S')$, and (2) it is possible for S' to have a conclusive reason for thinking that $\sim[P(S')\rightarrow Q(S')]$.

As has just been seen, all type I excluders are also type II excluders. Clearly, the converse is not the case. For example, "There are red lights shining on x" is a reason for thinking that $\sim(x$ looks red to $S\rightarrow x$ is red), but it is not a reason for thinking that x is not red. This suggests that we modify the definition of "type II excluder" to read "R is a type II excluder for the *prima facie* reason P for S to believe that Q, if and only if, R is a reason for S to believe that $\sim(P\rightarrow Q)$ but not a reason for S to believe that $\sim Q$." This has the effect of making the two categories of excluders mutually exclusive.

The fact that all excluders can be regarded as reasons for disbelieving the conditional $(P\rightarrow Q)$ gives us another way of looking at *prima facie* reasons. We can say that there are certain conditionals, such as the belief that if something looks red to me then it is red, that I am *prima facie justified in believing*. Or more precisely, there are certain conditionals for which it is necessarily true that (if one understands them) one is justified in believing them unless he thinks he has some reason for thinking they are false. We can say of such conditionals that there is a "logical presumption" in favor of believing them. This is a very important feature of *prima facie* reasons. Philosophers

have often been puzzled by how a person can know on the basis of perception that something is red without first ascertaining that there are no colored lights shining on the object, that he is not hallucinating or under the influence of drugs, that he is not hypnotized, etc. Analogously, philosophers have been puzzled by how a person can justifiably draw an inductive conclusion without first ascertaining that the inductive sample is a fair sample, but in the last analysis there does not appear to be any way to do this short of ascertaining that the inductive conclusion is true. This is a general puzzlement arising out of a misunderstanding of how type II excluders function. If P is a *prima facie* reason for one to believe that Q, it cannot be required that in order to employ this *prima facie* reason one first knows that there is nothing which, if known, would constitute a type II excluder. To know this would require that a person know on the basis of some sort of independent reasons that the conditional $(P\rightarrow Q)$ is true and thus no type II excluders can arise. But without relying upon the logical connection between the *prima facie* reason and the conclusion, the only way a person could know this is if he somehow already knew that the conclusion was true without relying upon the *prima facie* reason. This would lead to skepticism because in many cases (such as induction and color judgments) the *prima facie* reason constitutes the only basic way of getting to know whether the conclusion is true. If we were not allowed to use the *prima facie* reasons connected with induction or perception without first justifying them (in which case they would be merely *a posteriori* reasons), then we would never be able to draw any contingent general conclusions or make any color judgments. Thus the solution to the above puzzlement is simply the following. If P is a *prima facie* reason for S to believe that Q, then S does not need a reason for thinking that $(P\rightarrow Q)$. Evidence (for example, evidence concerning hallucination, drugs, colored lights, etc.) is *only* relevant if it is evidence *against* the conditional. Evidence *for* the conditional is never required. To suppose otherwise leads to skepticism, but skepticism is false.

The concept of a *prima facie* reason is like the concept of a logical reason in that both are purely relational concepts. And just as was the case for logical reasons, a *prima facie* reason need not be a good reason. If the belief-that-P is a *prima facie* reason for S to believe that Q, it is only a good reason for S to believe that Q if S also justifiably believes that P, understands the statements that Q, and does not think he has any excluding reasons.

Prima facie reasons are logically good reasons. The question naturally arises whether *prima facie* reasons exhaust the class of logically good reasons. My conjecture is that this is true, but I can see no way to prove it.

VI. Conclusion

The preceding picture of epistemic justification can be summarized as follows. Justification proceeds in terms of good reasons. Good reasons can be classified as logical reasons and *a posteriori* reasons. One class of logical reasons—conclusive reasons—is quite familiar. However these reasons do not exhaust the class of logical reasons. The more or less traditional assumption that they do has been a contributing factor in making a number of traditional epistemological problems appear insoluble. This is because the epistemic connection between such things as perception and the material world, or inductive grounds and the inductive conclusion, or behavior and mental states, cannot be described in terms solely of conclusive reasons. We must recognize at least one additional class of logical reasons—*prima facie* reasons. To map out precisely what the logical reasons are that are involved in these concepts is a difficult task, but with the help of the concept of a *prima facie* reason it no longer appears to be an impossible task, and hence the traditional epistemological difficulties concerning these concepts dissolve.

The reason that the existence of *prima facie* reasons has not generally been recognized is probably due to the difficulty in explaining within a traditional framework how it is possible for there to be such reasons. Philosophers have traditionally thought of the meaning of a statement as being determined solely by its truth conditions—the necessary and sufficient conditions for it to be true. But this is a mistake. The meaning of a statement is, of course, uniquely determined by its truth conditions, but these truth conditions can generally only be stated in a trivial way—by repeating the statement itself or by giving some rather uninteresting paraphrase of it. A more informative account of the meaning of a statement can be given by saying under what circumstances one would be justified in thinking the statement true, i.e., by saying what counts as a good reason for accepting the statement.[5] On this picture there is no difficulty in understanding how *prima facie* reasons are possible—they are just one kind of good reason that can be involved in making up the meaning of a statement.

State University of New York at Buffalo

[5] For a more exhaustive discussion of this point see my article, "What is an Epistemological Problem?", *American Philosophical Quarterly*, vol. 5 (1968), pp. 183–190.

V

Do Appearances Exist?

JOHN KNOX, JR.

I

THE most frequent objection to the term "appearance" is not that it has no meaning or legitimate use, but rather that one can so easily misuse it so as to suggest that appearances exist. Of course, if by "appearances exist" one means only that objects appear, then few would disagree. But one might mean that there exist certain items describable as "appearances." And many of our contemporaries would hold either that there are no such items, or that to say that there are, or that there are not, is a misuse of language. As these philosophers would see the matter, to affirm the existence of such items is to fall back into a neo-medieval darkness from which we have only lately emerged. I should like to suggest, however, that appearances may be defined in such a way that their existence as distinct items no longer is an absurdity, and indeed follows logically from the fact that objects are perceived. After attempting to explain and to justify this position, I shall conclude with the point that although the traditional problem of the logical foundations of empirical knowledge may have to be reconceived, one may, in fact one must, accept as both legitimate and important the old question of how an object is related to its appearances.

II

Usually, appearances are referred to as "sense data." But just what is a sense datum—i.e., what is the definition of "sense datum"? The traditional sense-datum philosophers (those philosophers who believe that there are, in a factual sense, certain items describable as "sense data") have been insufficiently concerned with the problem of definition. Since in their view the *existence* of sense data is relatively if not entirely uncontroversial, they have been content by and large to give directions for picking them out from among other, more ordinary existents. In particular, they have given such directions so as to provide concrete application for the otherwise largely empty concept of direct or immediate awareness. Aside from errors of psychology which may result, one trouble with such an approach

AMERICAN PHILOSOPHICAL
QUARTERLY

MONOGRAPH SERIES

Edited by NICHOLAS RESCHER

STUDIES IN THE THEORY OF KNOWLEDGE

Essays by:

Norman Malcolm
W. Donald Oliver
Peter Unger
John L. Pollock
John Knox, Jr.
Alan R. White
Frederick Stoutland

Monograph No. 4

Oxford, 1970

PUBLISHED BY BASIL BLACKWELL
WITH THE COOPERATION OF THE UNIVERSITY OF PITTSBURGH