

Externalist Theories of Empirical Knowledge

LAURENCE BONJOUR

Of the many problems that would have to be solved by a satisfactory theory of empirical knowledge, perhaps the most central is a general structural problem which I shall call *the epistemic regress problem*: the problem of how to avoid an infinite and presumably vicious regress of justification in one's account of the justification of empirical beliefs. *Foundationalist* theories of empirical knowledge, as we shall see further below, attempt to avoid the regress by locating a class of empirical beliefs whose justification does not depend on that of other empirical beliefs. *Externalist theories*, the topic of the present paper, represent one species of foundationalism.

I

I begin with a brief look at the epistemic regress problem. The source of the problem is the requirement that beliefs that are to constitute knowledge must be *epistemically justified*. Such a requirement is of course an essential part of the "traditional" conception of knowledge as justified true belief, but it also figures in at least most of the revisions of that conception which have been inspired by the Gettier problem. Indeed, if this requirement is understood in a sufficiently generic way, as meaning roughly that the acceptance of the belief must be epistemically rational, that it must not be epistemically irresponsible, then it becomes hard to see how any adequate conception of knowledge can fail to include it.

How then are empirical beliefs epistemically justified? Certainly the most obvious way to *show* that such a belief is justified is by producing a justificatory argument in which the belief to be justified is shown to follow inferentially from some other (perhaps conjunctive) belief, which is thus offered as a reason for accepting it. Beliefs whose justification would, if made explicit, take this form may be said to be *inferentially justified*. (Of course, such a justificatory argument would usually be explicitly rehearsed only in the face of some specific problem or challenge. Notice

also that an inferentially justified belief need not have been *arrived at* through inference, though it often will have been.)

The important point about inferential justification, however, is that if the justificandum belief is to be genuinely justified by the proffered argument, then the belief that provides the premise of the argument must itself be justified in some fashion. This premise belief might of course itself be inferentially justified, but this would only raise a new issue of justification with respect to the premise(s) of this new justificatory argument, and so on, so that empirical knowledge is threatened by an infinite and seemingly vicious regress of epistemic justification, with a *thoroughgoing* skepticism as the eventual outcome. So long as each new step of justification is inferential, it appears that justification can never be completed, indeed can never really even get started, and hence that there is no justification and no knowledge. Thus the epistemic regress problem.

What is the eventual outcome of this regress? There are a variety of possibilities, but the majority of philosophers who have considered the problem have believed that the only outcome that does not lead more or less directly to skepticism is *foundationalism*: the view that the regress terminates by reaching empirical beliefs (a) that are genuinely justified, but (b) whose justification is not inferentially dependent on that of any further empirical belief(s), so that no further issue of empirical justification is thereby raised. These non-inferentially justified beliefs, or *basic beliefs* as I shall call them, are claimed to provide the foundation upon which the edifice of empirical knowledge rests. And the central argument for foundationalism is simply that all other possible outcomes of the regress lead inexorably to skepticism.¹

This argument has undeniable force. Nonetheless, the central concept of foundationalism, the concept of a basic belief, is itself by no means unproblematic. The fundamental question that must be answered by any acceptable version of foundationalism is: *how are basic beliefs possible?* How, that is, is it possible for there to be an empirical belief that is epistemically justified in a way completely independent of any believed premises that might provide reasons for accepting it? As Chisholm suggests, a basic belief seems to be in effect an epistemologically unmoved (or perhaps self-moved) mover. But such a status is surely no less paradoxical in epistemology than it is in theology.

This intuitive difficulty with the idea of a basic empirical belief may be elaborated by considering briefly the fundamental concept of epistemic justification. There are two points to be made. First, the idea of justification is generic, admitting in principle of many different species. Thus, for example, the acceptance of an empirical belief might be morally justified, or pragmatically justified, or justified in some still different sense. But a belief's being justified in one of these other senses will not satisfy the justification condition for knowledge. What knowledge requires is *epistemic* justification. And the distinguishing characteristic of this particular species of justification is, I submit, its internal relationship to the cognitive goal of *truth*. A cognitive act is epistemically justified, on this conception, only if and to the extent that is aimed at this goal—which means at a minimum that one accepts only beliefs that there is adequate reason to think are true.

Second, the concept of epistemic justification is fundamentally a normative concept. It has to do with what one has a duty or obligation to do, from an epistemic or intellectual standpoint. As Chisholm suggests, one's purely intellectual duty is to accept beliefs that are true, or likely to be true, and reject beliefs that are false, or likely to be false. To accept beliefs on some other basis is to violate one's epistemic duty—to be, one might say, *epistemically irresponsible*—even though such acceptance might be desirable or even mandatory from some other, non-epistemic standpoint.

Thus if basic beliefs are to provide a suitable foundation for empirical knowledge, if inference from them is to be the sole basis for the justification of other empirical beliefs, then that feature, whatever it may be, in virtue of which an empirical belief qualifies as basic, must also constitute an adequate reason for thinking that the belief is true. And now if we assume, plausibly enough, that the person for whom a belief is basic must *himself* possess the justification for that belief if *his* acceptance of it is to be epistemically rational or responsible, and thus apparently that he must believe *with justification* both (a) that the belief has the feature in question and (b) that beliefs having that feature are likely to be true, then we get the result that this belief is not basic after all, since its justification depends on that of these other beliefs. If this result is correct, then foundationalism is untenable as a solution to the regress problem.²

What strategies are available to the foundationalist for avoiding this objection? One possibility would be to grant that the believer must be in possession of the reason for thinking that his basic belief is true but hold that the believer's cognitive grasp of that reason does not involve further *beliefs*, which would then require justification, but instead cognitive states of a different and more rudimentary kind: *intuitions* or *immediate apprehensions*, which are somehow capable of conferring justification upon beliefs without themselves requiring justification. Some such view as this seems implicit in most traditional versions of foundationalism.³

My concern in the present paper, however, is with an alternative foundationalist strategy, one of comparatively recent innovation. One way, perhaps somewhat tendentious, to put this alternative approach is to say that according to it, though there must in a sense be a reason why a basic belief is likely to be true, the person for whom such a belief is basic need not have any cognitive grasp of this reason. On this view, the epistemic justification or reasonableness of a basic belief depends on the obtaining of an appropriate relation, generally causal or nomological in character, between the believer and the world. This relation, which is differently characterized by different versions of the view, is such as to make it either nomologically certain or else highly probable that the belief is true. It would thus provide, *for anyone who knew about it*, an undeniably excellent reason for accepting such a belief. But according to proponents of the view under discussion, the person for whom the belief is basic need not (and in general will not) have any cognitive grasp of any kind of this reason or of the relation that is the basis for it in order for this basic belief to be justified; all these matters may be entirely *external* to the person's subjective conception of the situation. Thus the justification of a basic belief need not involve any further beliefs (or other cognitive states) so that no further regress of justification is

generated. D. M. Armstrong calls this an "externalist" solution to the regress problem, and I shall adopt this label.

My purpose in this paper is to examine such externalist views. I am not concerned with problems of detail in formulating a view of this kind, though some of these will be mentioned in passing, but rather with the overall acceptability of an externalist solution to the regress problem and thus of an externalist version of foundationalism. I shall attempt to argue that externalism is not acceptable. But there is a methodological problem with respect to such an argument which must be faced at the outset, since it determines the basic approach of the paper.

When viewed from the general standpoint of the western epistemological tradition, externalism represents a very radical departure. It seems safe to say that until very recent times, no serious philosopher of knowledge would have dreamed of suggesting that a person's beliefs might be epistemically justified simply in virtue of facts or relations that were external to his subjective conception. Descartes, for example, would surely have been quite unimpressed by the suggestion that his problematic beliefs about the external world were justified if only they were in fact reliably related to the world—whether or not he had any reason for thinking this to be so. Clearly his conception, and that of generations of philosophers who followed, was that such a relation could play a justificatory role only if the believer possessed adequate reason for thinking that it obtained. Thus the suggestion embodied in externalism would have been regarded by most epistemologists as simply irrelevant to the main epistemological issue, so much so that the philosopher who suggested it would have been taken either to be hopelessly confused or to be simply changing the subject (as I note below, this may be what some externalists in fact intend to be doing). The problem, however, is that this very radicalism has the effect of insulating the externalist from any very direct refutation: any attempt at such a refutation is almost certain to appeal to premises that a thoroughgoing externalist would not accept. My solution to this threatened impasse will be to proceed on an intuitive level as far as possible. By considering a series of examples, I shall attempt to exhibit as clearly as possible the fundamental intuition about epistemic rationality that externalism seems to violate. Although this intuition may not constitute a conclusive objection to the view, it is enough, I believe, to shift the burden of proof decisively to the externalist. In the final section of the paper, I shall consider briefly whether he can discharge this burden.

II

Our first task will be the formulation of a clear and relatively adequate version of externalism. The recent epistemological literature contains a reasonably large number of externalist and quasi-externalist views. Some of these, however, are not clearly relevant to our present concerns, either because they are aimed primarily at the Gettier problem, so that their implications for a foundationalist solution of the regress problem are not made clear, or because they *seem*, on the surface at least, to involve a repudiation of the very conception of epistemic justification or reasonableness as

a requirement for knowledge. Views of the latter sort seem to me to be very difficult to take seriously; but if they are seriously intended, they would have the consequence that the regress problem, at least in the form discussed here, would simply not arise, so that there would be no need for any solution, foundationalist or otherwise. My immediate concern here is with versions of externalism that claim to *solve* the regress problem and thus that also claim that the acceptance of beliefs satisfying the externalist conditions is epistemically justified or rational or warranted. Only such an externalist position genuinely constitutes a version of foundationalism, and hence the more radical views, if any such are in fact seriously intended, may safely be left aside for the time being.

The most completely developed externalist view of the sort we are interested in is that of Armstrong, as presented in his book, *Belief, Truth and Knowledge*.⁴ Armstrong is explicitly concerned with the regress problem, though he formulates it in terms of knowledge rather than justification. And it seems reasonably clear that he wants to say that beliefs satisfying his externalist criterion are epistemically justified or rational, though he is not as explicit as one might like on this point.⁵ In what follows, I shall in any case assume such an interpretation of Armstrong and formulate his position accordingly.

Another version of externalism, which fairly closely resembles Armstrong's except for being limited to knowledge derived from visual perception, is offered by Dretske in *Seeing and Knowing*.⁶ Goldman, in several papers, also suggests views of an externalist sort,⁷ and the view that Alston calls "Simple Foundationalism" and claims to be the most defensible version of foundationalism seems to be essentially externalist in character.⁸ The most extreme version of externalism would be one that held that the external condition required for justification is simply the *truth* of the belief in question. Such a view could not be held in general, of course, without obliterating the distinction between knowledge and mere true belief, thereby turning every lucky guess into knowledge. But it might be held with respect to some more limited class of beliefs. Such a view is mentioned by Alston as one possible account of privileged access,⁹ and seems, surprisingly enough, to be advocated by Chisholm (though it is very hard to be sure that this is what Chisholm really means).¹⁰

Here I shall concentrate mainly on Armstrong's view. Like all externalists, Armstrong makes the acceptability of a basic belief depend on an external relation between the believer and his belief, on the one hand, and the world, on the other, specifically a law-like connection: "there must be a *law-like connection* between the state of affairs *Bap* [i.e., *a*'s believing that *p*] and the state of affairs which makes '*p*' true, such that, given *Bap*, it must be the case that *p*." [166] This is what Armstrong calls the "thermometer-model" of non-inferential knowledge: just as the readings of a reliable thermometer lawfully reflect the temperature, so one's basic beliefs lawfully reflect the states of affairs that make them true. A person whose beliefs satisfy this condition is in effect a reliable cognitive instrument; and it is, according to Armstrong, precisely in virtue of this reliability that these basic beliefs are justified.

Of course, not all thermometers are reliable, and even a reliable one may be accurate only under certain conditions. Similarly, it is not a requirement for the jus-

tification of a basic belief on Armstrong's view that all beliefs of that general kind or even all beliefs of that kind held by that particular believer be reliable. Thus the law linking the having of the belief with the state of affairs that makes it true will have to mention properties, including relational properties, of the believer beyond his merely having that belief. Incorporating this modification yields the following schematic formulation of the conditions under which a non-inferential belief is justified and therefore basic: a non-inferential belief is justified if and only if there is some property *H* of the believer, such that it is a law of nature that whenever a person satisfies *H* and has that belief, then the belief is true. [197]¹¹ Here *H* may be as complicated as one likes and may include facts about the believer's mental processes, sensory apparatus, environment, and so on. But presumably, though Armstrong does not mention this point, *H* is not to include anything that would entail the truth of the belief; such a logical connection would not count as a law of nature.

Armstrong adds several qualifications to this account, aimed at warding off various objections, of which I shall mention only two. First, the nomological connection between the belief and the state of affairs that makes it true is to be restricted to "that of *completely reliable sign* to thing signified." [182] What this is intended to exclude is the case where the belief itself *causes* the state of affairs that makes it true. In such a case, it seems intuitively that the belief is not a case of knowledge even though it satisfies the condition of complete reliability formulated above. Second, the property *H* of the believer which is involved in the law of nature must not be "too specific"; there must be a "real possibility" of a recurrence of the situation described by the law. What Armstrong is worried about here is the possibility of a "veridical hallucination," i.e., a case in which a hallucinatory belief happens to be correct. In such a case, if the state of affairs that makes the belief true happens to be part of the cause of the hallucination and if the believer and his environment are described in enough detail, it might turn out to be nomologically necessary that such a state of affairs obtain, simply because all alternative possible causes for the hallucinatory belief have been ruled out by the specificity of the description. Again, such a case intuitively should not count as a case of knowledge, but it would satisfy Armstrong's criterion in the absence of this additional stipulation. (Obviously this requirement of non-specificity or repeatability is extremely vague and seems in fact to be no more than an *ad hoc* solution to this problem; but I shall not pursue this issue here.)

There are various problems of detail, similar to those just discussed, which could be raised about Armstrong's view, but these have little relevance to the main theme of the present paper. Here I am concerned with the more fundamental issue of whether Armstrong's view, or any other externalist view of this general sort, is acceptable as a solution to the regress problem and the basis for a foundationalist account of empirical knowledge. When considered from this perspective, Armstrong's view seems at the very least to be in need of considerable refinement in the face of fairly obvious counterexamples. Thus our first task will be to develop some of these counterexamples and suggest modifications in the view accordingly. This discussion will also lead, however, to a fundamental intuitive objection to all forms of externalism.

III

Although it is formulated in more general terms, the main concern of an externalist view like Armstrong's is obviously those non-inferential beliefs which arise from ordinary sources like sense-perception and introspection. For it is, of course, these beliefs which will on any plausible foundationalist view provide the actual foundations of empirical knowledge. Nevertheless, cases involving sense-perception and introspection are not very suitable for an intuitive assessment of externalism, since one central issue between externalism and other foundationalist and non-foundationalist views is precisely whether in such cases a further basis for justification beyond the externalist one is typically present. Thus it will be useful to begin by considering the application of externalism to other possible cases of non-inferential knowledge, cases of a less familiar sort where it will be easier to stipulate in a way that will be effective on an intuitive level that only the externalist sort of justification is present. Specifically, in this section and the next, our focus will be on possible cases of clairvoyant knowledge. Clairvoyance, the alleged psychic power of perceiving or intuiting the existence and character of distant states of affairs without the aid of any sensory input, remains the subject of considerable scientific controversy. Although many would like to dismiss out of hand the very idea of such a cognitive power, there remains a certain amount of evidence in favor of its existence which it is difficult to entirely discount. But in any case, the actual existence of clairvoyance does not matter at all for present purposes, so long as it is conceded to represent a coherent possibility. For externalism, as a general philosophical account of the foundations of empirical knowledge, must of course apply to all possible modes of non-inferential empirical knowledge, and not just to those that in fact happen to be realized.

The intuitive difficulty with externalism that the following discussion is intended to delineate and develop is this: on the externalist view, a person may be ever so irrational and irresponsible in accepting a belief, when judged in light of his own subjective conception of the situation, and may still turn out to be epistemically justified, i.e., may still turn out to satisfy Armstrong's general criterion of reliability. This belief may in fact be reliable, even though the person has no reason for thinking that it is reliable—or even though he has good reason to think that it is not reliable. But such a person seems nonetheless to be thoroughly irresponsible from an epistemic standpoint in accepting such a belief, and hence not justified, contrary to externalism. The following cases may help bring out this problem more clearly.

Consider first the following case:

Case I. Samantha believes herself to have the power of clairvoyance, though she has no reasons for or against this belief. One day she comes to believe, for no apparent reason, that the President is in New York City. She maintains this belief, appealing to her alleged clairvoyant power, even though she is at the same time aware of a massive amount of apparently cogent evidence, consisting of news reports, press releases, allegedly live television pictures, etc., indicating that the President is at that time in Washington, D.C. Now the President is in fact in New York City, the evidence to the contrary being part of a massive of-

ficial hoax mounted in the face of an assassination threat. Moreover, Samantha does in fact have completely reliable clairvoyant power, under the conditions that were then satisfied, and her belief about the President did result from the operation of that power.

In this case, it is clear that Armstrong's criterion of reliability is satisfied. There will be some complicated description of Samantha, including the conditions then operative, from which it will follow, by the law describing her clairvoyant power, that her belief is true.¹² But it seems intuitively clear nevertheless that this is not a case of justified belief or of knowledge: Samantha is being thoroughly irrational and irresponsible in disregarding cogent evidence that the President is not in New York City on the basis of a clairvoyant power which she has no reason at all to think that she possesses; and this irrationality is not somehow canceled by the fact that she happens to be right. Thus, I submit, Samantha's irrationality and irresponsibility prevent her belief from being epistemically justified.

This case and others like it suggest the need for a further condition to supplement Armstrong's original one: not only must it be true that there is a law-like connection between a person's belief and the state of affairs that makes it true, such that given the belief, the state of affairs cannot fail to obtain, but it must also be true that the person in question does not possess cogent reasons for thinking that the belief in question is false. For, as this case seems to show, the possession of such reasons renders the acceptance of the belief irrational in a way that cannot be overridden by a purely externalist justification.

Nor is this the end of the difficulty for Armstrong. Suppose that the clairvoyant believer, instead of having evidence against the particular belief in question, has evidence against his possession of such a cognitive power, as in the following case:

Case II. Casper believes himself to have the power of clairvoyance, though he has no reasons for this belief. He maintains his belief despite the fact that on the numerous occasions on which he has attempted to confirm one of his allegedly clairvoyant beliefs, it has always turned out apparently to be false. One day Casper comes to believe, for no apparent reason, that the President is in New York City, and he maintains this belief, appealing to his alleged clairvoyant power. Now in fact the President is in New York City; and Casper does, under the conditions that were then satisfied, have completely reliable clairvoyant power, from which this belief in fact resulted. The apparent falsity of his other clairvoyant beliefs was due in some cases to his being in the wrong conditions for the operation of his power and in other cases to deception and misinformation.

Is Casper justified in believing that the President is in New York City, so that he then knows that this is the case? According to Armstrong's account, even with the modification just suggested, we must apparently say that the belief is justified and hence a case of knowledge: the reliability condition is satisfied, and Casper possesses no reason for thinking that the President is not in New York City. But this result still seems mistaken. Casper is being quite irrational and irresponsible from an epistemic

standpoint in disregarding evidence that his beliefs of this sort are not reliable and should not be trusted. And for this reason, the belief in question is not justified.

In the foregoing case, Casper possessed good reasons for thinking that he did not possess the sort of cognitive ability that he believed himself to possess. But the result would be the same, I believe, if someone instead possessed good reasons for thinking that *in general* there could be no such cognitive ability, as in the following case:

Case III. Maud believes herself to have the power of clairvoyance, though she has no reasons for this belief. She maintains her belief despite being inundated by her embarrassed friends and relatives with massive quantities of apparently cogent scientific evidence that no such power is possible. One day Maud comes to believe, for no apparent reason, that the President is in New York City, and she maintains this belief, despite the lack of any independent evidence, appealing to her alleged clairvoyant power. Now in fact the President is in New York City, and Maud does, under the conditions then satisfied, have completely reliable clairvoyant power. Moreover, her belief about the President did result from the operation of that power.

Again, Armstrong's criterion of reliability seems to be satisfied. But it also seems to me that Maud, like Casper, is not justified in her belief about the President and does not have knowledge. Maud has excellent reasons for thinking that no cognitive power such as she believes herself to possess is possible, and it is irrational and irresponsible of her to maintain her belief in that power in the face of that evidence and to continue to accept and maintain beliefs on this dubious basis.

Cases like these two suggest the need for a further modification of Armstrong's account: in addition to the law-like connection between belief and truth and the absence of any reasons against the particular belief in question, it must also be the case that the believer in question has no cogent reasons, either relative to his own case or in general, for thinking that such a law-like connection does *not* exist, i.e., that beliefs of that kind are not reliable.

IV

So far the modifications suggested for Armstrong's criterion are consistent with the basic thrust of externalism as a response to the regress problem. What emerges is in fact a significantly more plausible externalist position. But these cases and the modifications made in response to them also suggest an important moral which leads to a basic intuitive objection to externalism: external or objective reliability is not enough to offset subjective irrationality. If the acceptance of a belief is seriously unreasonable or unwarranted from the believer's own standpoint, then the mere fact that unbeknownst to the believer its existence in those circumstances lawfully guarantees its truth will not suffice to render the belief epistemically justified and thereby an instance of knowledge. So far we have been concerned only with situations in which the believer's subjective irrationality took the form of ignoring positive grounds in his possession for questioning either that specific belief or beliefs arrived at in that

way. But now we must ask whether even in a case where these positive reasons for a charge of irrationality are not present, the acceptance of a belief where only an externalist justification is available cannot still be said to be subjectively irrational in a sense that rules out its being epistemically justified.

We may begin by considering one further case of clairvoyance, in which Armstrong's criterion with all the suggested modifications is satisfied:

Case IV. Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.

Is Norman epistemically justified in believing that the President is in New York City, so that his belief is an instance of knowledge? According to the modified externalist position, we must apparently say that he is. But is this the right result? Are there not still sufficient grounds for a charge of subjective irrationality to prevent Norman's being epistemically justified?

One thing that might seem relevant to this issue, which I have deliberately omitted from the specification of the case, is whether Norman *believes* himself to have clairvoyant power, even though he has no justification for such a belief. Let us consider both possibilities. Suppose, first, that Norman does have such a belief and that it contributes to his acceptance of his original belief about the President's whereabouts in the sense that were Norman to become convinced that he did not have this power, he would also cease to accept the belief about the President.¹³ But is it not obviously irrational, from an epistemic standpoint, for Norman to hold such a belief when he has no reasons at all for thinking that it is true or even for thinking that such a power is possible? This belief about his clairvoyance fails after all to possess even an externalist justification. And if we say that the belief about his clairvoyance is epistemically irrational and unjustified, must we not say the same thing about the belief about the President which *ex hypothesi* depends upon it?¹⁴

A possible response to this challenge would be to add one further condition to our modified externalist position, *viz.*, that the believer not even *believe* that the law-like connection in question obtains, since such a belief will not in general be justified (or at least that his continued acceptance of the particular belief that is at issue not depend on his acceptance of such a general belief). In our present case, this would mean that Norman must not believe that he has the power of clairvoyance (or at least that his acceptance of the belief about the President's whereabouts not depend on his having such a general belief). But if this specification is added to the case, it now becomes more than a little puzzling to understand what Norman thinks is going on. From his standpoint, there is apparently no way in which he *could* know the President's whereabouts. Why then does he continue to maintain the belief that

the President is in New York City? Why is not the mere fact that there is no way, as far as he knows or believes, for him to have obtained this information a sufficient reason for classifying this belief as an unfounded hunch and ceasing to accept it? And if Norman does not do this, is he not thereby being epistemically irrational and irresponsible?

For these reasons, I submit, Norman's acceptance of the belief about the President's whereabouts is epistemically irrational and irresponsible, and thereby unjustified, whether or not he believes himself to have clairvoyant power, so long as he has no justification for such a belief. Part of one's epistemic duty is to reflect critically upon one's beliefs, and such critical reflection precludes believing things to which one has, to one's knowledge, no reliable means of epistemic access.¹⁵

We are now face-to-face with the fundamental—and seemingly obvious—intuitive problem with externalism: *why* should the mere fact that such an external relation obtains mean that Norman's belief is epistemically justified, when the relation in question is entirely outside his ken? As remarked earlier, it is clear that one who knew that Armstrong's criterion was satisfied would be in a position to construct a simple and quite cogent justifying argument for the belief in question: if Norman has property *H* (being a completely reliable clairvoyant under the existing conditions and arriving at the belief on that basis), then he holds the belief in question only if it is true; Norman does have property *H* and does hold the belief in question; therefore, the belief is true. But Norman himself is by stipulation not in a position to employ this argument, and it is unclear why the mere fact that it is, so to speak, potentially available in the situation should justify *his* acceptance of the belief. Precisely what generates the regress problem in the first place, after all, is the requirement that for a belief to be justified for a particular person, not only is it necessary that there be true premises somehow available in the situation which could in principle provide a basis for a justification, but also that the believer in question know or at least justifiably believe some such set of premises and thus be in a position to employ the corresponding argument. The externalist position seems to amount merely to waiving this general requirement in a certain class of cases, and the question is why this should be acceptable in these cases when it is not acceptable generally. (If it were acceptable generally, then it seems likely that *any* true belief would be justified, unless some severe requirement is imposed as to how immediately available such premises must be. But any such requirement seems utterly arbitrary, once the natural one of actual access by the believer is abandoned.) Thus externalism looks like a purely *ad hoc* solution to the epistemic regress problem.

One reason why externalism may seem initially plausible is that if the external relation in question genuinely obtains, then Norman will in fact not go wrong in accepting the belief, and it is, *in a sense*, not an accident that this is so. But how is this supposed to justify Norman's belief? From his subjective perspective, it *is* an accident that the belief is true. Of course, it would not be an accident from the standpoint of our hypothetical external observer who knows all the relevant facts and laws. Such an observer, having constructed the justifying argument sketched above, would be thereby in a position to justify *his own* acceptance of the belief. Thus

Norman, as Armstrong's thermometer image suggests, could serve as a useful epistemic instrument for such an observer, a kind of cognitive thermometer; and it is to this fact, as we have seen, that Armstrong appeals in arguing that a belief like Norman's can be correctly said to be reasonable or justifiable. [183] But none of this seems in fact to justify Norman's *own* acceptance of the belief, for Norman, unlike the hypothetical external observer, has no reason at all for thinking that the belief is true. And the suggestion here is that the rationality or justifiability of Norman's belief should be judged from Norman's own perspective, rather than from one that is unavailable to him.¹⁶

This basic objection to externalism seems to me to be intuitively compelling. But it is sufficiently close to being simply a statement of what the externalist wants to deny to make it helpful to buttress it a bit by appealing to some related intuitions.

First, we may consider an analogy with moral philosophy. The same conflict between perspectives which we have seen to arise in the process of epistemic assessment can also arise with regard to the moral assessment of a person's action: the agent's subjective conception of what he is doing may differ dramatically from that which would in principle be available to an external observer who had access to facts about the situation that are beyond the agent's ken. And now we can imagine an approximate moral analogue of externalism which would hold that the moral justifiability of an agent's action was, in certain cases at least, properly to be determined from the external perspective, entirely irrespective of the agent's own conception of the situation.

Consider first the moral analogue of Armstrong's original, unmodified version of externalism. If we assume, purely for the sake of simplicity, a utilitarian moral theory, such a view would say that an action might on occasion be morally justified simply in virtue of the fact that in the situation then obtaining, it would as a matter of objective fact lead to the best overall consequences—even if the agent planned and anticipated that it would lead to a very different, perhaps extremely undesirable, consequences. But such a view seems plainly mistaken. There is no doubt a point to the objective, external assessment: we can say correctly that it turns out to be objectively a good thing that the agent performed the action. But this is not at all inconsistent with saying that his action was morally unjustified and reprehensible, given his subjective conception of the likely consequences.

Thus our envisaged moral externalism must at least be modified in a way that parallels the modifications earlier suggested for epistemological externalism. Without attempting to make the analogy exact, it will suffice for our present purposes to add to the original requirement for moral justification, viz., that the action will in fact lead to the best overall consequences, the further condition that the agent not believe or intend that it lead to undesirable consequences. Since it is also, of course, not required by moral externalism that the agent believe that the action will lead to good consequences, the sort of case we are now considering is one in which an agent acts in a way that will in fact produce the best overall consequences, but has *no belief at all* about the likely consequences of his action. Although such an agent is no doubt preferable to one who acts in the belief that his action will lead to undesirable consequences, surely he is not morally justified in what he does. On the contrary,

he is being highly irresponsible, from a moral standpoint, in performing the action in the absence of any evaluation of what will result from it. His moral duty, from our assumed utilitarian standpoint, is to do what will lead to the best consequences, but this duty is not satisfied by the fact that he produces this result willy-nilly, without any idea that he is doing so.¹⁷ And similarly, the fact that a given sort of belief is objectively reliable, and thus that accepting it is in fact conducive to arriving at the truth, need not prevent our judging that the epistemic agent who accepts it without any inkling that this is the case violates his epistemic duty and is epistemically irresponsible and unjustified in doing so.

Second, we may appeal to the connection between knowledge and rational action. Suppose that Norman, in addition to the clairvoyant belief described earlier, also believes that the Attorney-General is in Chicago. This latter belief, however, is not a clairvoyant belief but is based upon ordinary empirical evidence in Norman's possession, evidence strong enough to give the belief some fairly high degree of reasonableness, but *not* strong enough to satisfy the requirement for knowledge.¹⁸ Suppose further that Norman finds himself in a situation where he is forced to bet a very large amount, perhaps even his life or the life of someone else, on the whereabouts of either the President or the Attorney-General. Given his epistemic situation as described, which bet is it more reasonable for him to make? It seems relatively clear that it is more reasonable for him to bet the Attorney-General is in Chicago than to bet that the President is in New York City. But then we have the paradoxical result that from the externalist standpoint it is more rational to act on a merely reasonable belief than to act on one that is adequately justified to qualify as knowledge (and which in fact *is* knowledge). It is very hard to see how this could be so. If greater epistemic reasonableness does not carry with it greater reasonableness of action, then it becomes most difficult to see why it should be sought in the first place. (Of course, the externalist could simply bite the bullet and insist that it is in fact more reasonable for Norman to bet on the President's whereabouts than the Attorney-General's, but such a view seems very implausible.)

I have been attempting in this section to articulate the fundamental intuition about epistemic rationality, and rationality generally, that externalism seems to violate. This intuition the externalist would of course reject, and thus my discussion does not constitute a refutation of the externalist position on its own ground. Nevertheless it seems to me to have sufficient intuitive force at least to place the burden of proof squarely on the externalist. In the final section of the paper, I shall consider briefly some of the responses that seem to be available to him.

V

One possible defense for the externalist in the face of the foregoing intuitive objection would be to narrow his position by restricting it to those commonsensical varieties of non-inferential knowledge which are his primary concern, viz., sense-perception and introspection, thereby rendering the cases set forth above strictly irrelevant. Such a move seems, however, utterly *ad hoc*. Admittedly it is more dif-

difficult to construct intuitively compelling counterexamples involving sense-perception and introspection, mainly because our intuitions that beliefs of those kinds are in fact warranted in *some way* or other are very strong. But this does nothing to establish that the externalist account of their warrant is the correct one. Thus unless the externalist can give some positive account of why the same conclusion that seems to hold for non-standard cases like clairvoyance does not also hold for sense-perception and introspection, this narrowing of his position seems to do him no good.

If the externalist cannot escape the force of the objection in this way, can he perhaps balance it with positive arguments in favor of his position? Many attempts to argue for externalism are in effect arguments by elimination and depend on the claim that alternative accounts of empirical knowledge are unacceptable, either because they cannot solve the regress problem or for some other reason. Most such arguments, depending as they do on a detailed consideration of the alternatives, are beyond the scope of the present paper. But one such argument depends only on very general features of the competing positions and thus can usefully be considered here.

The basic factual premise of this argument is that in very many cases that are commonsensically instances of justified belief and of knowledge, there seem to be no justifying factors explicitly present beyond those appealed to by the externalist. An ordinary person in such a case may have no idea at all of the character of his immediate experience, of the coherence of his system of beliefs, etc., and yet may still have knowledge. Alternative theories, so the argument goes, may describe correctly cases of knowledge involving a knower who is extremely reflective and sophisticated, but they are obviously too demanding and too grandiose when applied to these more ordinary cases. In these cases, *only* the externalist condition is satisfied, and this shows that no more than that is necessary for justification and for knowledge, though more might still be epistemically desirable.

Although the precise extent to which it holds could be disputed, in the main this factual premise must be simply conceded. Any non-externalist account of empirical knowledge that has any plausibility will impose standards for justification which very many beliefs that seem commonsensically to be cases of knowledge fail to meet in any full and explicit fashion. And thus on such a view, such beliefs will not *strictly speaking* be instances of adequate justification and of knowledge. But it does not follow that externalism must be correct. This would follow only with the addition of the premise that the judgments of common sense in this area are sacrosanct, that any departure from them is enough to demonstrate that a theory of knowledge is inadequate. But such a premise seems entirely too strong. There seems in fact to be no basis for more than a reasonably strong presumption in favor of the correctness of common sense, but one which is still quite defeasible. And what it would take to defeat this presumption depends in part on how great a departure from common sense is being advocated. Thus, although it would take very strong grounds to justify a very strong form of skepticism, not nearly so much would be required to make acceptable the view that what common sense regards as cases of justification and of knowledge are in fact only rough approximations to an epistemic ideal which *strictly speaking* they do not satisfy.

Of course, a really adequate reply to the externalist would have to spell out in some detail the precise way in which such beliefs really do approximately satisfy some acceptable alternative standard, a task which obviously cannot be attempted here. But even without such elaboration, it seems reasonable to conclude that this argument in favor of externalism fails to carry very much weight as it stands and would require serious buttressing in order to give it any chance of offsetting the intuitive objection to externalism: either the advocacy and defense of a quite strong presumption in favor of common sense, or a detailed showing that alternative theories cannot in fact grant to the cases favored by common sense even the status of approximations to justification and to knowledge.

The other pro-externalist argument I want to consider does not depend in any important way on consideration of alternative positions. This argument is hinted at by Armstrong [185–88], among others, but I know of no place where it is developed very explicitly. Its basic claim is that only an externalist theory can handle a certain version of the lottery paradox.

The lottery paradox is standardly formulated as a problem confronting accounts of inductive logic that contain a rule of acceptance or detachment, but we shall be concerned here with a somewhat modified version. This version arises when we ask how much or what degree of epistemic justification is required for a belief to qualify as knowledge, given that the other necessary conditions for knowledge are satisfied. Given the intimate connection, discussed earlier, between epistemic justification and likelihood of truth, it seems initially reasonable to take likelihood or probability of truth as a measure of the degree of epistemic justification, and thus to interpret the foregoing question as asking how likely or probable it must be, relative to the justification of one's belief, that the belief be true, in order for that belief to satisfy the justification requirement for knowledge. Most historical theories of knowledge tended to answer that knowledge requires *certainty* of truth, relative to one's justification. But more recent epistemological views have tended to reject this answer, for familiar reasons, and to hold instead that knowledge requires only a reasonably high likelihood of truth. And now, if this high likelihood of truth is interpreted in the obvious way as meaning that, relative to one's justification, the numerical probability that one's belief is true must equal or exceed some fixed value, the lottery paradox at once rears its head.

Suppose, for example, that we decide that a belief is adequately justified to satisfy the requirement for knowledge if the probability of its truth, relative to its justification, is 0.99 or greater. Imagine now that a lottery is to be held, about which we know the following facts: exactly 100 tickets have been sold, the drawing will indeed be held, it will be a fair drawing, and there will be only one winning ticket. Consider now each of the 100 propositions of the form:

Ticket number n will lose

where n is replaced by the number of one of the tickets. Since there are 100 tickets and only one winner, the probability of each such proposition is 0.99; and hence if we believe each of them, our individual beliefs will be adequately justified to satisfy

the requirement for knowledge. And then, given only the seemingly reasonable assumptions, first, that if one has adequate justification for believing each of a set of propositions, one also has adequate justification for believing the conjunction of the members of the set, and, second, that if one has adequate justification for believing a proposition, one also has adequate justification for believing any further proposition entailed by the first proposition, it follows that we are adequately justified in believing that no ticket will win, contradicting our other information.

Clearly this is a mistaken result, but how is it to be avoided? In the first place, it will plainly do no good to simply increase the level of numerical probability required for adequate justification. For no matter how high it is raised, short of certainty, it will obviously be possible to duplicate the paradoxical result by simply choosing a large enough lottery. Nor do the standard responses to the lottery paradox, whatever their merits may be in dealing with other versions of the paradox, seem to be of much help here. Most of them are ruled out simply by insisting that we do know that empirical propositions are true, not merely that they are probable, and that such knowledge is not in general relative to particular contexts of inquiry. This leaves only the possibility of avoiding the paradoxical result by rejecting the two assumptions stated in the preceding paragraph. But this would be extremely implausible—involving in effect a denial that one may always justifiably deduce conclusions from one's putative knowledge—and in any case would still leave the intuitively unacceptable result that one could on this basis come to know separately the 99 true propositions about various tickets losing (though not of course the false one). In fact, it seems intuitively clear that I do not *know* any of these propositions to be true: if I own one of the tickets, I do not know that it will lose, even if in fact it will, and would not know no matter how large the total number of tickets might be.

At this stage, it may seem that the only way to avoid the paradox is to return to the traditional idea that any degree of probability or likelihood of truth less than certainty is insufficient for knowledge, that only certainty, relative to one's justification, will suffice. The standard objection to such a view is that it seems to lead at once to the skeptical conclusion that we have little or no empirical knowledge. For it seems quite clear that there are no empirical beliefs, with the possible and extremely problematic exception of beliefs about one's own mental states, for which we have justification adequate to exclude all possibility of error. Such a solution seems as bad as the original problem.

It is at this point that externalism may seem to offer a way out. For an externalist position allows one to hold that the justification of an empirical belief must make it certain that the belief is true, while still escaping the clutches of skepticism. This is so precisely because the externalist justification need not be within the cognitive grasp of the believer or indeed of anyone. It need only be true that there is *some* description of the believer, however complex and practically unknowable it may be, which, together with *some* true law of nature, ensures the truth of the belief. Thus, e.g., my perceptual belief that there is a cup on my desk is not certain, on any view, relative to the evidence or justification that is in my possession; I might be hallucinating or there might be an evil demon who is deceiving me. But it seems reasonable

to suppose that if the belief is indeed true, then there is *some* external description of me and my situation and *some* true law of nature, relative to which the truth of the belief is guaranteed, and if so it would satisfy the requirement for knowledge.

In some ways, this is a neat and appealing solution to the paradox. Nonetheless, it seems doubtful that it is ultimately satisfactory. In the first place, there is surely something intuitively fishy about solving the problem by appealing to an in-principle guarantee of truth which will almost certainly in practice be available to no one. A second problem, which cannot be elaborated here, is that insisting on this sort of solution seems likely to create insuperable difficulties for knowledge of general and theoretical propositions. But in any case, the externalist solution seems to yield intuitively incorrect results in certain kinds of cases. A look at one of these may also suggest the beginnings of a more satisfactory solution.

Consider then the following case:

Case V. Agatha, seated at her desk, believes herself to be perceiving a cup on the desk. She also knows, however, that she is one of a group of 100 people who have been selected for a philosophical experiment by a Cartesian evil demon. The conditions have been so arranged that all 100 will at this particular time seem to themselves to be perceiving a cup upon their respective desks, with no significant differences in the subjective character of their respective experiences. But in fact, though 99 of the people will be perceiving a cup in the normal way, the last one will be caused by the demon to have a complete hallucination (including perceptual conditions, etc.) of a non-existent cup. Agatha knows all this, but she does not have any further information as to whether she is the one who is hallucinating, though as it happens she is not.

Is Agatha epistemically justified in her belief that there is a cup on the desk and does she know this to be so? According to the externalist view, we must say that she is justified and does know. For there is, we may assume, an external description of Agatha and her situation relative to which it is nomologically certain that her belief is true. (Indeed, according to Armstrong's original, unmodified view, she would be justified and would know even if she also knew instead that 99 of the 100 persons were being deceived by the demon, so long as she was in fact the odd one who was perceiving normally.) But this result is, I submit, intuitively mistaken. If Agatha knows that she is perceiving a cup, then she also knows that she is not the one who is being deceived. But she does not know this, for reasons that parallel those operative in the lottery case.

Is there then no way out of the paradox? The foregoing case and others like it seem to me to suggest the following approach to at least the present version of the paradox, though I can offer only an exceedingly brief sketch here. Intuitively, what the lottery case and the case of Agatha have in common is the presence of a large number of relevantly similar, alternative possibilities, all individually very unlikely, but such that the person in question *knows* that at least one of them will in fact be realized. In such a case, since there is no relevant way of distinguishing among these possibilities, the person cannot believe with adequate justification and *a fortiori* can-

not know that any particular possibility will not be realized, even though the probability that it will not be realized may be made as high as one likes by simply increasing the total number of possibilities. Such cases do show that high probability is not by itself enough to satisfy the justification condition for knowledge. They do not show, however, that certainty is required instead. For what rules out knowledge in such a case is not merely the fact that the probability of truth is less than certainty but also the fact that the person *knows* that at least one of these highly probable propositions is false. It is a necessary condition for justification and for knowledge that this not be so. But there are many cases in which a person's justification for a belief fails to make it certain that the belief is true, but in which the person also does not know that some possible situation in which the belief would be false is one of a set of relevantly similar, alternative possibilities, at least one of which will definitely be realized. And in such a case, the lottery paradox provides no reason to think that the person does not know.¹⁹

An example may help to make this point clear. Consider again my apparent perception of the cup on my desk. I think that I do in fact know that there is a cup there. But the justification that is in my possession surely does not make it certain that my belief is true. Thus, for example, it seems to be possible, relative to my subjective justification, that I am being deceived by an evil demon, who is causing me to have a hallucinatory experience of the cup, together with accompanying conditions of perception. But it does not follow from this that I do not know that there is a cup on the desk, because it does not follow and I do not know that there is some class of relevantly similar cases in at least one of which a person is in fact deceived by such a demon. Although it is only probable and not certain that there is no demon, it is still possible for all I *know* that never in the history of the universe, past, present, or future, is there a case in which someone in a relevantly similar perceptual situation is actually deceived by such a demon. And, as far as I can see, the same thing is true of all the other ways in which it is possible that my belief might be mistaken. If this is so, then the lottery paradox provides no obstacle to my knowledge in this case.²⁰

This response to the lottery paradox seems to me to be on the right track. It must be conceded, however, that it is in considerable need of further development and may turn out to have problems of its own. But that is a subject for another paper.²¹

There is one other sort of response, mentioned briefly above, which the externalist might make to the sorts of criticisms developed in this paper. I want to remark on it briefly, though a full-scale discussion is impossible here. In the end it may be possible to make intuitive sense of externalism only by construing the externalist as simply abandoning the traditional idea of epistemic justification or rationality and along with it anything resembling the traditional conception of knowledge. I have already mentioned that this may be precisely what the proponents of externalism intend to be doing, though most of them are anything but clear on this point.²²

Against an externalist position that seriously adopts such a gambit, the criticisms developed in the present paper are of course entirely ineffective. If the externalist does not want even to claim that beliefs satisfying his conditions are epistemically justified or reasonable, then it is obviously no objection that they seem in some

cases to be quite unjustified and unreasonable. But, as already noted, such a view, though it may possess some other sort of appeal, constitutes a solution to the epistemic regress problem or to any problem arising out of the traditional conception of knowledge only in the radical and relatively uninteresting sense that to reject that conception entirely is also, of course, to reject any problems arising out of it. Such "solutions" would seem to be available for any philosophical problem at all, but it is hard to see why they should be taken seriously.

Notes

1. For a fuller discussion of the regress argument, including a discussion of other possible outcomes of the regress, see my paper "Can Empirical Knowledge Have a Foundation?" *American Philosophical Quarterly* 15 (1978):1-13. That paper also contains a brief anticipation of the present discussion of externalism.

2. It could, of course, still be claimed that the belief in question was *empirically* basic, so long as both the needed justifying premises were justifiable on an *a priori* basis. But this would mean that it was an *a priori* truth that a particular empirical belief was likely to be true. In the present paper, I shall simply assume, without further discussion, that this seemingly unlikely state of affairs does not in fact obtain.

3. For criticism of this view, see the paper cited in note 1.

4. D. M. Armstrong, *Belief, Truth and Knowledge* (London, Press, 1973). Bracketed references in the text will be to the pages of this book.

5. The clearest passages are at p. 183, where Armstrong says that a belief satisfying his externalist condition, though not "based on reasons," nevertheless "might be said to be reasonable (justifiable), because it is a sign, a completely reliable sign, that the situation believed to exist does in fact exist"; and at p. 189, where he suggests that the satisfaction of a slightly weaker condition, though it does not yield knowledge, may still yield rational belief. There is no reason to think that any species of rationality or reasonableness other than the epistemic is at issue in either of these passages. But though these passages seem to me to adequately support my interpretation of Armstrong, the strongest support may well derive simply from the fact that he at no point *disavows* a claim of epistemic rationality. (See also the parenthetical remark in the middle of p. 77.)

6. Fred I. Dretske, *Seeing and Knowing* (London, 1969), chap. III. Dretske also differs from Armstrong in requiring in effect that the would-be knower also believe that the externalist condition is satisfied, but not of course that this belief be justified.

7. Goldman does this most clearly in "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73 (1976):771-91; and in "What is Justified Belief?" forthcoming. See also "A Causal Theory of Knowing," *Journal of Philosophy* 64 (1967):355-72, though this last paper is more concerned with the Gettier problem than with a general account of the standards of epistemic justification.

8. William P. Alston, "Two Types of Foundationalism," *Journal of Philosophy* 73 (1976): 165-85; see especially p. 168.

9. Alston, "Varieties of Privileged Access," in Roderick Chisholm and Robert Swartz, *Empirical Knowledge* (Englewood Cliffs, N.J., 1973), pp. 396-99. Alston's term for this species of privileged access is "truth-sufficiency."

10. See Chisholm, *Theory of Knowledge*, 2nd ed. (Englewood Cliffs, N.J., 1977), p. 22, where Chisholm offers the following definition of the concept of a state of affairs being *self-presenting*:

b is *self-presenting* for *S* at *t* = Df *b* occurs at *t*; and necessarily, if *b* occurs at *t* then *b* is evident [i.e., justified] for *S* at *t*.

Despite the overtones of the term "self-presentation," nothing in this passage seems to require that believer have any sort of immediate awareness of the state in question; all that is required is that it actually occur, i.e., that his belief be true. On the other hand, Chisholm also, in the section immediately preceding this definition, quotes with approval a passage from Leibniz which appeals to the idea of "direct awareness" and of the absence of mediation "between the understanding and its objects," thus suggesting the non-externalist variety of foundationalism (pp.20-21).

11. Armstrong actually formulates the criterion as a criterion of knowledge, rather than merely of justification; the satisfaction of the belief condition is built into the criterion and this, with the satisfaction of the indicated justification condition, entails that the truth condition is satisfied.

12. This assumes that clairvoyant beliefs are caused in some distinctive way, so that an appropriately complete description of Samantha will rule out the possibility that the belief is a mere hunch and will connect appropriately with the law governing her clairvoyance.

13. This further supposition does not prevent the belief about the President's whereabouts from being non-inferential, since it is not in any useful sense Norman's reason for accepting that specific belief.

14. This is the basic objection to Dretske's version of externalism, mentioned above. Dretske's condition requires that one have an analogously unjustified (though true) belief about the reliability of one's perceptual belief.

15. The only apparent answer here would be to claim that the reasonable presumption is in favor of one's having such reliable means of access, unless there is good reason to the contrary. But it is hard to see why such a presumption should be thought reasonable.

16. Mark Pastin, in a critical study of Armstrong, has suggested that ascriptions of knowledge depend on the epistemic situation of the ascriber rather than on that of the ascribee at this point, so that I am correct in ascribing knowledge to Norman so long as I know that his belief is reliable (and hence also that the other conditions of knowledge are satisfied), even if Norman does not. But I can see no very convincing rationale for this claim. See Pastin, "Knowledge and Reliability: A Study of D. M. Armstrong's *Belief, Truth and Knowledge*," *Metaphilosophy* 9 (1978):150-62. Notice further that if the epistemic regress problem is in general to be dealt with along externalist lines, then my knowledge that Norman's belief is reliable would depend on the epistemic situation of a further external observer, who ascribes knowledge to me. And similarly for the knowledge of that observer, etc., *ad infinitum*. I do not know whether this regress of external observers is vicious, but it seems clearly to deprive the appeal to such an observer of any value as a practical criterion.

17. Of course there are cases in which one must act, even though one has no adequate knowledge of the likely consequences; and one might attempt to defend epistemic externalism by arguing that in epistemic contexts the analogous situation *always* obtains. But there are several problems with such a response. First, to simply assume that this is always so seems to be question-begging, and the externalist can argue for this claim only by refuting all alternatives to his position. Second, notice that in ethical contexts this situation usually, perhaps always, obtains only when not acting will lead definitely to bad consequences, not just to the failure to obtain good ones; and there seems to be no parallel to this in the epistemic case. Third, and most important, the justification for one's action in such a case would depend not on the external fact, if it is a fact, that the action leads to good consequences, but simply on the fact that one could do no better, given the unfortunate state of one's knowledge; thus this position would not be genuinely a version of moral externalism, and analogously for the epistemic case.

18. I am assuming here, following Chisholm, that knowledge requires a degree of justification stronger than that required to make a belief merely reasonable.

19. I do not, alas, have any real account to offer here of the notion of *relevant similarity*. Roughly, the idea is that two possibilities are relevantly similar if there is no known difference between them that has a bearing on the likelihood that they will be realized. But this will not quite do. For consider a lottery case in which there are two tickets bearing each even number and only one for each odd number. Intuitively, it seems to me, this difference does not prevent all the tickets, odd and even, from being relevantly similar, despite the fact that it is twice as likely that an even ticket will be drawn.

20. But if this account is correct, I may still fail to know in many other cases in which common sense would say fairly strongly that I do. E.g., do I know that my house has not burned down since I left it this morning? Ordinarily we are inclined to say that we do know such things. But if it is true, as it might well be, that I also know that of the class of houses relevantly similar to mine, at least one will burn down at some point, then I do not, on the present account, *know* that my house has not burned down, however improbable such a catastrophe may be. (On the other hand, knowledge would not be ruled out by the present principle simply because I knew that certain specific similar houses, *other than mine*, have in the past burned down or even that they will in the future burn down. For I know, *ex hypothesi*, that my house is not one of those. The force of the principle depends on my knowing that at least one possibility *which might for all I know be the one I am interested in* will be realized, not just on descriptively similar possibilities being realized.)

21. This response to the lottery paradox derives in part from discussions with C. Anthony Anderson.

22. The clearest example of such a position is in Goldman's paper "Discrimination and Perceptual Knowledge," cited above, where he rejects what he calls "Cartesian-style justification" as a requirement for perceptual knowledge, in favor of an externalist account. He goes on to remark, however, that one could use the term "justification" in such a way that satisfaction of his externalist conditions "counts as justification," though a kind of justification "entirely different from the sort of justification demanded by Cartesianism" (p. 790). What is unclear is whether this is supposed to be a purely verbal possibility, which would then be of little interest, or whether it is supposed to connect with something like the concept of epistemic rationality explicated in section I. Thus it is uncertain whether Goldman means to repudiate the whole idea of epistemic rationality, or only some more limited view such as the doctrine of the given (reference to which provides his only explanation of what he means by "Cartesianism" in epistemology).