

*The Structure of  
Empirical Knowledge*

---

LAURENCE BONJOUR

*Harvard University Press*

CAMBRIDGE, MASSACHUSETTS, AND LONDON, ENGLAND

1985

---

crucial cases will not) have any cognitive grasp of this reason, or of the relation that is the basis of it, in order for his belief to be justified; all of this may be entirely *external* to his 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 and the fundamental foundationalist problem is neatly solved.'

The recent epistemological literature contains a number of externalist and quasi-externalist views. Some of these, however, are not clearly relevant to my 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 very difficult to take seriously; but they would in any case have the consequence that the regress problem in the form with which we are concerned would simply not arise, so that there would be no need for any solution, foundationalist or otherwise. My concern here is with the versions of externalism that attempt to *solve* the regress problem by claiming 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 be safely ignored here.

Perhaps the most completely developed externalist view of this sort is Armstrong's, as presented in his book *Belief, Truth, and Knowledge*.<sup>2</sup> Armstrong is explicitly concerned with the regress problem, though he formulates it in terms of knowledge rather than justification. And it is reasonably clear that he wants to say that beliefs which satisfy his externalist criterion are epistemically justified or rational, though he is not as explicit as one might like on this point.<sup>3</sup> In what follows I will in any case assume such an interpretation of Armstrong and formulate his position accordingly.<sup>4</sup>

Like all externalist foundationalists, Armstrong makes the justification 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 lawlike connection: "there must be a *law-like connection* between the state of affairs *Bap* [such as *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 also one's basic beliefs lawfully reflect the states of affairs which make them true. A person whose beliefs satisfy

## 3 Externalist Versions of Foundationalism

### 3.1 The basic idea of externalism

In the last chapter I argued that empirical foundationalism faces a serious and fundamental problem, which is that any foundationalist view must somehow manage the feat of (a) avoiding any requirement that the believer have further justified empirical beliefs to provide reasons for thinking that his allegedly basic empirical beliefs are true (which would destroy their status as basic), while still (b) maintaining in some way the essential connection between justification and truth.

The externalist response to this problem amounts to the claim that although there must indeed exist a reason why a basic empirical belief is likely to be true (or even, in some versions, guaranteed to be true), the person for whom the belief is basic need not himself have any cognitive grasp at all of this reason (thus rejecting premise (3) of the antifoundationalist argument). Instead, it is claimed, the epistemic justification or reasonableness of a basic empirical belief derives from the obtaining of an appropriate relation, generally construed as causal or nomological in character, between the believer and the world. This relation, which is differently characterized by different versions of externalism, is such as to make it either nomologically certain or highly probable that the belief is true. It would thus provide, *for anyone who knew about it*, an undeniably excellent reason for accepting the belief. But according to externalism, the person for whom the belief is basic need not (and in the

this condition is in effect a reliable cognitive instrument; and it is, according to Armstrong, precisely in virtue of this reliability that his 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 justification of a basic belief according to Armstrong's view that all beliefs of that general kind or even all beliefs of that kind held by that particular believer be reliable. The law linking the having of the belief with the state of affairs which makes it true will have to mention properties of the believer, including relational properties, 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 noninferential 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).<sup>5</sup> 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.

Armstrong adds several qualifications to this account, aimed at warding off various objections, of which I will mention only one. The nomological connection between the belief and the state of affairs which makes it true is to be restricted to "that of *completely reliable sign* to thing specified" (182). What this is intended to exclude is the case where the belief itself *causes* the state of affairs which makes it true. In such a case, it seems intuitively clear that the belief is not knowledge even though it satisfies the condition of complete reliability formulated above.

There are various problems of detail, similar to the one just discussed, which could be raised about Armstrong's view, but these have little relevance to the main concerns of this book. I am concerned with the more fundamental issue of whether Armstrong's view, or any externalist position of this general sort, is acceptable as a solution to the regress problem and as the basis for a foundationalist account of empirical knowledge. I will attempt to argue that externalism is not acceptable. But there is a serious methodological problem with respect to such an argument which must be faced at the outset, since it determines the basic approach of the rest of this chapter (which differs substantially from the balance of the book).

When viewed from the general standpoint of the Western epistemological tradition, externalism represents a quite substantial 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 merely in virtue of facts or relations that are 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 caused, 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 himself possessed adequate reasons for thinking that the relation obtained. Thus the suggestion embodied in externalism would have been regarded 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 already noted, this *may* be what some externalists intend to be doing).

My own conviction is that this reaction is in fact correct, that externalism (like a number of other distinctively "analytic" solutions to classical philosophical problems) reflects an inadequate appreciation of the problem at which it is aimed. But the problem is how to argue for this view — assuming that one is unwilling simply to dismiss externalism out of hand. For this very radicalism has the effect of insulating the externalist from any 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 as far as possible on an intuitive level. By considering a series of examples, I will attempt first to suggest some needed refinements in Armstrong's view and eventually to exhibit clearly the fundamental intuition concerning epistemic rationality that externalism violates. Although this intuition may not constitute a conclusive objection to the view, it is enough, I submit, to place the burden of proof squarely on the externalist. In the later sections of the chapter, I will consider whether there is any way in which he can discharge this burden.

### 3.2 Some counter-examples to Armstrong's view

Although it is formulated in more general terms, the main concern of an externalist view like Armstrong's is clearly those noninferential beliefs which arise from familiar sources like sense-perception and introspection, for it is these beliefs which will on any plausible foundationalist view provide the actual foundations of empirical knowledge. But cases involving sense-perception and introspection are nevertheless not very suitable for an intuitive assessment of externalism, since one central issue between externalism and other foundationalist and nonfoundationalist 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 noninferential knowledge: cases of a less familiar sort for which it will be easier to stipulate in a way which will be effective on an intuitive level that *only* the externalist sort of justification is present. Specifically, in this section and the next, my 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 controversy; although it is hard not to be skeptical about such an exotic cognitive power, the alleged evidence in favor of its existence is difficult to discount entirely. 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 noninferential empirical knowledge, not just to those which happen to be realized.

The intuitive difficulty with externalism which the following discussion is intended to suggest and develop is this: according to the externalist view, a person may be highly 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 according to Armstrong's criterion. His belief may in fact be reliable, even though he has no reason for thinking it is reliable—or even 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 in fact justified. The following cases may help to bring out this problem more clearly.

Consider first this case:

**Case 1.** Samatha 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, and so on, 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 official hoax mounted in the face of an assassination threat. Moreover, Samantha does in fact have completely reliable clairvoyant power under the conditions which 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 presumably quite complicated description of Samantha, including the conditions then operative, from which it will follow via the law describing her clairvoyant power that her belief is true.<sup>6</sup> But it seems clear nevertheless that this is not a case of justified belief or of knowledge. Samantha is being thoroughly irrational and irresponsible in disregarding the 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 the case that there is a lawlike connection between a person's belief and the state of affairs which makes it true such that given the belief, the state of affairs cannot fail to obtain, but it must also be the case that the person does not possess cogent reasons for thinking that the belief 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 2.** 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 when 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 which were 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 or misinformation.

Is Casper justified in believing that the President is in New York City, and does he therefore know 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 as an individual did not possess the sort of cognitive ability which he believed himself to possess. But the result would be the same, I submit, if someone instead possessed good reasons for thinking that *in general* there could be no such cognitive ability, as in the following case:

*Case 3.* 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 is 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 lawlike connection between belief and truth and the absence of 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 situation or in general, for thinking that such a lawlike connection does *not* exist, that is, that beliefs of that kind are not reliable.

### 3.3 A basic objection to externalism

Up to this point the suggestive modifications of 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 him 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 consists in ignoring positive grounds in his possession for questioning either that specific belief or beliefs arrived at in that general 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 which rules out its being epistemically justified.

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

*Case 4.* Norman, under certain conditions which 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 conclusion? Aren't there still sufficient grounds for a charge of subjective irrationality to prevent Norman from being epistemically justified?

One thing that may 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 the 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.' 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 have even an externalist justification. And if we say that the belief about his clairvoyance is epistemically unjustified, must we not say the same thing about the belief about the President which *ex hypothesi* depends upon it?<sup>8</sup>

A possible response to this challenge would be to add one further condition to our modified externalist position, namely, that the believer not even *believe* that the lawlike connection in question obtains (or at least that his continued acceptance of the particular belief which is at issue not depend on his acceptance of such a belief), since such a belief cannot in general be justified. In case 4, 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 must not depend on his having such a belief). But if this specification is added to the case, it becomes quite difficult to understand what Norman himself 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 isn't the mere fact that there is no way, as far as he knows, 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, isn't he thereby being epistemically irrational and irresponsible?

Thus, 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.<sup>9</sup>

We are now face-to-face with the fundamental—and 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 I

noted 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 that the President is in New York City: 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. Such an external observer, having constructed this justifying argument, would be thereby in a position to justify *his own* acceptance of a belief with the same content. Thus Norman, as Armstrong's own 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 is *ex hypothesi* 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 it is necessary not only that there be true premises or reasons somehow available in the situation that 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 or reasons and thus be *himself* in a position to offer the corresponding justification. 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 were imposed as to how available such premises must be; and 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: it would not be an accident from the standpoint of our hypothetical external observer who knows all the relevant facts and laws. But how is this supposed to justify Norman's belief? From his subjective perspective, it *is* an accident 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 which is unavailable to him."

This basic objection to externalism seems to me 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, consider an analogy with moral philosophy. The same conflict between perspectives which has been 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 of an external observer who has access to facts about the situation which 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 is, 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 be morally justified simply in virtue of the fact that in the situation then obtaining it would lead as a matter of objective fact to the best overall consequences—even though the agent himself planned and anticipated that it would lead to very different, extremely undesirable consequences. But such a view seems 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 did what he did, his bad intentions notwithstanding. But this is not at all inconsistent with saying that his action was morally unjustified and reprehensible in light of his subjective conception of the likely consequences.

Thus our envisioned moral externalism must at least be modified in a way which 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 requirement for moral justification just envisioned (that the action will in fact lead to the best overall consequences) the further condition that the agent not believe or intend that it will lead to undesirable consequences. Since it is also, of course, not required by moral externalism that he believe that the action will lead to the best consequences, the 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 himself *no belief at all* about the likely consequences of his action. But while 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 doing what he does in the absence of any conception of what will result. 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.<sup>11</sup> 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, consider the connection between knowledge and rational action. Suppose that Norman, in addition to having the clairvoyant belief described earlier, also believes that the Attorney General is in Chicago. This latter belief, however, is not a clairvoyant belief but rather is based on ordinary empirical evidence in Norman's possession, evidence strong enough to give the belief a fairly high degree of reasonableness, but *not* strong enough to satisfy the requirement for knowledge. 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 clear that it is more reasonable for him to bet that 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 which 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.

I have been attempting in this section to articulate the fundamental intuition concerning epistemic rationality, and rationality generally, that externalism seems to violate. This intuition would of course be rejected by the externalist, and thus my discussion does not constitute a refutation of externalism on its own ground. Nevertheless, it seems to me to have sufficient intuitive force at the very least to shift the burden of proof strongly to the externalist. In the rest of this chapter I will consider what responses are available to him.

### 3.4 Some externalist rejoinders

There are several ways in which an externalist might respond to the foregoing argument. Perhaps the least attractive is to simply stand his ground: adhere stubbornly to Armstrong's original position and claim that even in cases 1–3 the beliefs in question are epistemically justified in spite of the apparent subjective irrationality involved. One way in which one might try to make such a position plausible is to argue that the idea of epistemic irresponsibility that is the basis of the intuitive argument above is simply inapplicable to beliefs of the sort in question, or at least to their more ordinary perceptual and introspective analogues, because such beliefs are essentially involuntary in character (an involuntariness that is obscured by the use of the term “accept,” which misleadingly suggests deliberate action). How can I be irresponsible, the argument would go, in doing something which I cannot help doing?<sup>12</sup>

There are, however, two difficulties with this response. First, while it is true that beliefs of this sort, or indeed perhaps of any sort, are not voluntary in any simple way, it is wrong to regard them as involuntary to the degree which this view requires. While one may not be able to decide simply not to accept such a belief, one can, especially over an extended period of time, “bracket” the belief: refuse to take it seriously, to draw any conclusions from it, to act upon it, and so on. This is essentially what happens in the case of known perceptual illusions, compelling but unsubstantiated hunches, and recognized prejudices, for example, and there is no reason why the same treatment could not be accorded to other sorts of beliefs which are superficially involuntary. Second, even if the response in question were sufficient to show that the holding of such beliefs is not irresponsible or irrational in itself, it would have no tendency at all to establish that it is reasonable to employ such beliefs as premises for the derivation of other beliefs in the way that foundationalism requires.

A more promising line for the externalist is to accept the result advocated here for cases 1–3, together with the suggested modifications in Armstrong's position, but attempt to avoid the extension of that result to the crucial case 4, claiming that in that case the purely externalist basis for justification does suffice to render Norman's belief epistemically justified. Such a reply, if successfully defended, would save all that is really vital to the externalist position. But how can case 4 be successfully prised apart from the earlier ones? What the externalist needs at this point is a different account of *why* the beliefs in cases 1–3 are not justified, an account which does not invoke the notion of subjective irrationality and hence does not extend readily to case 4.

Such an alternative account is suggested by Alvin Goldman.<sup>13</sup> Having offered an externalist position basically similar to Armstrong's (though differing in requiring only that the process which produces justified beliefs generally produce true ones, in contrast to Armstrong's stronger requirement that it always do so), Goldman suggests the following case as a possible counter-example (I have restated it slightly):

*Case 5.* Jones is told on reliable authority that a certain class of his memory beliefs are entirely mistaken. His parents fabricate a wholly false story that he suffered from amnesia when he was age seven and later developed completely false “pseudo-memories” of the forgotten period in his life. Though Jones has excellent reasons to trust his parents, he persists in believing the ostensible memories from the period in question.<sup>14</sup>

This case obviously parallels our earlier case 2. The beliefs in question, assuming that they in fact result from normal processes of memory, would be justified according to Goldman's initial position and even, if a few additional details are filled in appropriately, according to Armstrong's stronger position.

Goldman, however, agrees with the intuition that such beliefs are not justified and is thus forced to modify his view. After considering and rejecting some alternative possibilities, he arrives at the following revised condition for when a belief is justified:

If S's belief in *p* at *t* results from a reliable cognitive process, and there is no reliable . . . process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing *p* at *t*, then S's belief in *p* at *t* is justified.<sup>15</sup>

According to this condition, as interpreted by Goldman, we can say that the beliefs in case 5, and also those in cases 1–3, are not justified; in each of these cases there is a reliable cognitive process, which Goldman describes as “the proper use of evidence,” that would if used have led the persons in question not to accept those beliefs. Whereas in case 4 there is no such reliable process available to Norman that if employed would have led to an alteration in his belief. Thus Goldman's revised position provides a different analysis of what has gone wrong in cases like 1–3 and 5 (an available and reliable cognitive process has not been used), an analysis which does not appear to generalize to case 4, thus leaving the central externalist position untouched.

But of course the mere existence of an alternative account does not establish that this account is correct. We have a set of cases (1–3, 5) in



which we may take it as agreed that external reliability does not suffice for justification, together with two alternative ways of extrapolating from these cases to a more general principle: one that appeals to subjective irrationality, which would thus also class Norman's belief in case 4 as unjustified; and one that appeals to the presence of an available and reliable alternative cognitive process, which would not have such a result. How are we to decide between these two extrapolations? I have already said in the preceding section all that I can in favor of the former one. Thus the present issue is what can be said for, and against, Goldman's alternative.

Unfortunately, however, Goldman's positive rationale for his view is quite difficult to make out clearly. Indeed, what little he does say, if given a sensitive reading, seems to tell against his view and in favor of the subjective rationality alternative: "the proper use of evidence would be an instance of a . . . reliable process. So what we can say about Jones is that he *fails* to use a certain . . . reliable process that he could and should have used. . . he failed to do something which, epistemically, he should have done . . . The justificational status of a belief is not only a function of the cognitive processes *actually* employed in producing it; it is also a function of processes that could *and should* be employed."<sup>16</sup> The obvious problem here is how to interpret the suggestion that Jones *should* have used the alternative cognitive process in question. On the surface this seems to be an appeal to the idea of subjective rationality and as such would favor the alternative position. I have been unable to arrive at any alternative construal of the passage (except one which would make the idea that Jones should use the other method merely a reiteration of the fact that it is reliable); and certainly to omit the phrases in question would make this passage intuitively much less satisfactory as an account of case 5.

Moreover, it seems relatively easy to think of cases in which Goldman's condition yields the wrong result. These are cases in which there is an alternative cognitive process available that is in fact reliable, but that the person in question has no reason to think is reliable. Thus consider the following case:

*Case 6.* Cecil is a historian and is concerned to answer a certain historical question. After spending a large amount of time on his research and consulting all of the available sources and documents, he accumulates a massive and apparently conclusive quantity of evidence in favor of a certain answer to his question. He proceeds to accept that answer, which is in fact correct. At the same time, however, Cecil happens to have in his possession a certain crystal ball; and in fact the answers given by this

crystal ball are extremely, but not perfectly, reliable with regard to the sort of subject matter in question, though Cecil hasn't the slightest reason to suspect this (he also has no reason to think that crystal balls are not reliable). Moreover, the crystal ball would, if consulted, have given a different answer to the question at issue (one of its rare mistakes); and Cecil, if he had consulted the crystal ball and accorded to its answer the degree of evidential weight corresponding to its degree of reliability, would have been led to accept neither answer to his question.

Is Cecil epistemically justified in accepting the belief in question? According to Goldman's condition we must say that he is not, for consulting the crystal ball (and taking its answers seriously) is an alternative cognitive process which is both available and reliable, but which would, if employed, have led to his not accepting the belief. But this answer seems mistaken, so long as Cecil has no reason to think that the alternative process is reliable. Thus Goldman's revised position is not acceptable, and this general sort of externalist response to the objection raised here thus does not succeed.

There is, however, a further and rather different externalist response which must be considered. The intuitive argument against externalism was formulated, for reasons already discussed, in terms of an admittedly rather anomalous variety of noninferential knowledge—one which is certainly quite possible, as far as we can tell, but whose empirical credentials are nevertheless at present dubious at best. But the externalist's primary concern is not such nonstandard cases but rather those familiar varieties of noninferential knowledge that can reasonably be assumed to provide the actual foundation upon which empirical knowledge rests, if it rests on a foundation at all: noninferential knowledge deriving from sense-perception and introspection. The application of the view to clairvoyance and similar cases is quite inessential to the main thrust of the externalist position. For this reason, an obvious and initially appealing response for the externalist would be simply to pull in his horns, abandon the unnecessarily general form of his view discussed above, and advocate it only as restricted to the range of cases which are his main interest. I will call such a view, which it will not be necessary to formulate exactly, *restricted externalism*.

Can such a retrenchment save externalism from the intuitive force of the objection offered above? It must be conceded that such a restricted externalism initially seems more plausible than the more comprehensive version, but it is very doubtful whether it is really any better off. Though the anti-externalist argument was formulated in terms of clairvoyance, the conception of epistemic rationality which it puts forward—of such

rationality as essentially dependent on the believer's own subjective conception of his epistemic situation—was and is intended to be perfectly general in its application. Having in effect accepted that argument as applied to nonstandard cases like clairvoyance (for otherwise why restrict his Position?), the restricted externalist must explain clearly why it does not apply equally well to the more familiar cases with which he is concerned. If mere external reliability is not sufficient to epistemically justify a clairvoyant belief, *why* does it somehow become adequate in the case of a sensory belief or an introspective one? What is the difference between the two sorts of cases?

It is crucial at this point to see clearly that the restricted externalist cannot evade this issue by simply relying on the greater intuitive appeal of his limited position. Though this appeal has already been conceded, it is quite possible that it derives covertly from factors to which the externalist may not legitimately appeal. Thus one difference between cases of clairvoyance and cases of sense-perception or introspection might be that cases of the latter sort involve immediately given or intuited subjective experience which somehow provides a basis for justification but which is sufficiently tacit in its operation as to yield the mistaken impression that only externalist factors are at work. Such an appeal to subjective experience would represent a version of the doctrine of the given, which I will consider (and, in fact, reject) in Chapter 4. A second possibility, of more interest in connection with this book since it roughly approximates the positive view I will offer in Part II, is that the difference between beliefs deriving from sources like clairvoyance, on the one hand, and sensory or introspective beliefs, on the other, depends on the believer's being epistemically justified in thinking that beliefs of the latter sort are in fact generally reliable (though again this dependence is tacit enough to be easily overlooked). According to this view, if such beliefs were in fact reliable but the believer in question did not know this at least implicitly, then they would not be justified. Now on neither of these accounts is the externalist basis for justification in fact sufficient for justification; the intuitive impression that it is sufficient is based on overlooking crucial, though inconspicuous, aspects of the situation. And if this is so, the initial intuitive appeal of restricted externalism is spurious. Of course, these alternative accounts have not been shown to be correct. But neither have they been shown by the externalist to be incorrect. And failing such a showing, the only way for the restricted externalist to make it plausible that such factors are not or need not be at work is to provide an alternative account of the difference between cases like clairvoyance and his favored cases of sense-perception and introspection, an account which vindicates the initial intuition that restricted externalism is more

acceptable, that the argument given above for cases like clairvoyance does not apply here.

Is any such account available? While the possibility of one cannot, of course, be ruled out entirely, none has in fact been offered, and there is no apparent reason for thinking that there is any important asymmetry in this respect between clairvoyance, on the one hand, and sense-perception and introspection, on the other. Consider first cases of sense-perception. For any particular sense it would be possible to formulate cases parallel to those formulated for clairvoyance involving positive grounds for distrusting a specific belief (cases 1–3). And the intuitive result as regards justification would, I submit, be the same so long as nonexternalist justifying factors were clearly excluded, though this is harder to see because the reliability of the senses is normally taken so completely for granted (perhaps the easiest way to see the point is to envisage the discovery of a new sense, one not known in advance to be reliable). Thus subjective rationality is relevant to the justification of sensory beliefs to at least this extent. But once this degree of parallelism between the two cases is admitted, it is hard to see why the further argument of section 3.3 is not also applicable to sense-perception.

The case of introspection is somewhat more complicated. There is no reason to think that a person who holds an introspective belief could not have cogent grounds for doubting the correctness of that particular belief, thus yielding an analogue of case 1. Such evidence might be behavioral in character or might result from the use of some sort of brain-scanning device. But it is a consequence of the positive view of empirical justification I will offer in Part II that there could be no introspective analogues of cases 2 and 3.<sup>17</sup> The reason is that on that account *all* empirical justification depends on the premise, presumably derived from introspection, that one has a certain system of beliefs; one could not have cogent empirical reasons for thinking that one's introspective beliefs were generally unreliable without undercutting this essential premise and thus empirical justification generally—including the justification of those very reasons themselves. The conclusion that introspection is unreliable in general might be *true*, but one could never be justified in believing it. Thus at least part of the reason for saying that justification depends on subjective rationality is not available in the case of introspection, if the position I will eventually defend here is correct.

Does this mean that an externalist account of introspection can escape the general line of objection developed in this chapter? I do not believe that it does. In the first place, there is still the analogue of case 1 to show that subjective rationality is essential to justification. Second, the unavailability of analogues of cases 2 and 3 *for the reason just dis-*

*cussed* has no real tendency to show that a purely external basis of justification is somehow more adequate here than in other cases; it makes it harder to argue against introspective externalism directly, but it does nothing to undermine the claim that introspection is epistemically analogous to other cases like clairvoyance and sense-perception. Third, the other considerations advanced in section 3.3 can still be applied directly to this species of externalism, and these seem to have considerable force on their own. For these reasons it does not seem that restricted externalism is a defensible retreat for the externalist.<sup>18</sup>

### 3.5 Arguments in favor of externalism

If the externalist cannot escape by retreating to restricted externalism, can he perhaps balance the objections 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 this chapter. But there is one which 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 many cases which are commonsensically instances of justified belief and of knowledge, there seem to be no justifying factors 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, or of whatever other basis of justification a nonexternalist position may appeal to, and yet may still have knowledge. Alternative theories, so the argument goes, may perhaps describe correctly cases of knowledge involving a knower who is extremely reflective and sophisticated, but they are obviously too demanding and grandiose when applied to these more mundane cases. In these cases *only* the externalist condition is satisfied, and this shows that no more than that is really necessary for justification, and for knowledge, though more might still be in some sense epistemically desirable.

Though the precise extent to which it holds could be disputed, in the main the initial factual premise of this argument must simply be conceded. Any nonexternalist account of empirical knowledge that has any plausibility will impose standards for justification that many com-

monsensical cases of knowledge will fail to meet in any full and explicit way. 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 as to which of our beliefs qualify as knowledge are sacrosanct, that any serious departure from them is enough to demonstrate that a theory of knowledge is inadequate. But, as already discussed in connection with Chisholm's "problem of the criterion" (in section 1.3), such a premise seems entirely too strong. There seems in fact to be no basis for more than a quite defeasible *presumption* (if indeed even that) in favor of the correctness of common sense. And what it would take to defeat this presumption depends in part on how great a departure from common sense is being advocated. Thus while it would take very strong grounds to justify a strong form of skepticism which claims that the beliefs which common sense regards as knowledge have no significant positive epistemic status at all, not nearly so much would be required to make acceptable the view that these beliefs are in fact only rough approximations to an epistemic ideal which *strictly speaking* they do not satisfy.<sup>19</sup>

Of course, a really adequate reply to this argument would have to spell out in some detail the precise way in which such beliefs really do approximately satisfy the standards in question, and I will attempt to do this in the development of the positive account advocated below. But even without such elaboration, it seems reasonable to conclude that this argument in favor of externalism fails to have much weight as it stands. To give it any chance of offsetting the intuitive objection to externalism developed earlier would require either the advocacy and defense of a much stronger presumption in favor of common sense than seems at all obviously correct or else a showing that alternative theories cannot in fact grant to the cases favored by common sense even the status of approximations to justifications and to knowledge. And until such buttressing is forthcoming, this argument may safely be set aside.

The other pro-externalist argument that I want to consider is one which does not depend in any important way on consideration of alternative positions. This argument is hinted at by Armstrong (185–188), among others, but I know of no place where it is developed very explicitly. Its basic claim is that only an externalist theory can solve 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 will be concerned here with a somewhat modified version. This version arises when we ask what degree of epistemic jus-

tification is required for a belief to qualify as knowledge, assuming that the other necessary conditions for knowledge are satisfied. Given the intimate connection, discussed above, 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 adequate-justification requirement for knowledge. Many historical theories of knowledge have answered that knowledge requires *certainty* of truth relative to one's justification. But more recent views have tended to reject this answer, on the grounds that it leads inevitably to an automatic and uninteresting skepticism, and to hold instead that knowledge requires only a reasonably high likelihood or probability of truth. And if such a 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 .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 .99; and hence if I believe each of them, my 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 those propositions; 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 I am adequately justified in believing that no ticket will win, contradicting my other information.

Clearly this is a mistaken result, but how is it to be avoided? It will obviously do no good simply to increase the level of numerical probability required for adequate justification, for no matter how high it is raised, short of certainty, it will be possible to duplicate the paradoxical result merely by choosing a sufficiently large lottery. Nor do the standard responses to the lottery paradox, whatever their merits may be in dealing with other versions, seem to be of much help here. Most of them may

be ruled out simply by insisting that we do know that some empirical propositions are true, not merely that they are probable, and that such knowledge is not in general relative to particular contexts of inquiry. Of the standard solutions, this leaves only the possibility of avoiding the paradoxical result by rejecting the two assumptions stated in the preceding paragraph. But such a rejection 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 repugnant result that one could on this basis come to know separately the *yy* true propositions about various tickets losing (though not, of course, the false one). In fact, however, 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 this is so 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 of truth will suffice, a solution which threatens to lead at once to skepticism. It is at this point that externalism might appear to help. For an externalist position would allow one to hold, following Armstrong, that the factors which justify an empirical belief must make it nomologically 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 the case that there is *some* description of the believer, however complex and practically unknowable it may be, which together with *some* true law of nature, again perhaps practically unknowable, ensures the truth of the belief. Thus, for example, my perceptual belief that there is a cup on my desk is not certain, on any view, relative to the evidence or justification which 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 there is *some* external description of me and my situation and *some* true law of nature relative to which the belief is certain; and *if so*, it satisfies the externalist requirements for knowledge.

I doubt, however, whether this superficially neat solution to the paradox is ultimately satisfactory. In the first place, there is surely something intuitively fishy about solving the problem by appeal to a theoretical guarantee of truth which will almost certainly be in practice available to *no* one. A second problem 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, such as the following one:

*Case 7.* 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 by a Cartesian demon for a philosophical experiment. The conditions have been so arranged by the demon that all 100 will at this particular time seem to themselves to be perceiving a cup on their respective desks with no significant differences in the subjective character of their respective experiences. But in fact, though 99 of these 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 perceived perceptual conditions) of a nonexistent 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 to a degree adequate for knowledge in her belief that there is a cup on the desk? 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 version of externalism, she would be justified and would know even if she also knew that 99 of the 100 persons, instead of only one, 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 suggest, 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 exactly parallel to those which prevent a person in the original lottery case from knowing that his ticket will lose. Thus externalism fails to provide a correct solution to this version of the paradox.<sup>21</sup>

There is one other sort of response, mentioned briefly above, which the externalist might want to make to the sorts of criticisms developed in this chapter. I want to remark on it briefly, though a full-scale discussion is beyond the scope of this work. 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 some proponents of externalism intend to be doing, though most of them are anything but clear on this point.<sup>22</sup> Against an externalist

position which seriously adopts such a gambit, the criticisms developed in this chapter are, of course, entirely ineffective. If the externalist does not mean to claim that beliefs which satisfy his conditions are epistemically justified or reasonable, then it is obviously no objection to his view that they seem in some cases to be quite unjustified and unreasonable. But such a view, though it may be in some other way attractive or useful, 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 to reject any problems arising out of it. In this book I will confine myself to less radical solutions.<sup>23</sup>