

Level-Confusions in Epistemology

WILLIAM P. ALSTON

Uncovering confusions in each other's work is a favorite, almost, one sometimes suspects, the sole, occupation of contemporary American philosophers. I am surely not the only member of this class who has to resist temptations to spend a disproportionate amount of time on such activities. After all, it is so much easier than presenting and defending substantive theses. And it is a lot of fun. Like the rest of fallen humanity, I resist temptation only part of the time, and this is, I fear, the other part. In this paper I will be engaged in uncovering what I take to be some fundamental and pervasive confusions in contemporary epistemology. However, in this instance I have more solid reasons than usual for spending time in confusion spotting. I do think that epistemology is one area in which the practitioners, even (or perhaps especially) the most significant ones, have fallen into certain confusions that have profoundly influenced their systematic constructions. Hence by revealing those confusions one can make an important contribution to the development of epistemology with relatively little effort. At least that is my claim for what I am doing in this paper. You can form your own judgement as to whether it is correct.

The confusions to which I will be calling your attention all involve sloughing over the distinction between epistemic levels, proceeding as if what is true of a proposition, belief, or epistemic state of affairs on one level is *ipso facto* true of a correlated proposition, belief, or epistemic state of affairs on another. The levels I have in mind are those built up by the introduction and iteration of epistemic or pistic operators: 'know that', 'believe that', 'is justified in believing that', and so on. Thus if we begin with any proposition, p , we can build a structure of epistemic levels by using various epistemic operators.

- I. p
- S believes that p
- S believes that S believes that p
-

II. *p*

S is justified in believing that *p*
 S is justified in believing that S is justified in believing that *p*

.....

III. *p*

S knows that *p*
 S knows that S knows that *p*

.....

We can also have "mixed" items. *S knows that p* can give rise to the higher-level *S believes that S knows that p* or the equally higher-level *S is justified in believing that S knows that p*. My purposes in this paper do not require me to develop precise criteria for determining the relative levels of any two such items. The confusions we will be disclosing are all between items that are obviously on different levels.

I

My first example concerns the concept of immediate (direct) justification. The contrast between mediate (indirect) and immediate (direct) justification can be most simply and most fundamentally stated as follows.

- (1) To say that a belief is meditately justified is to say that what justifies it includes some other justified beliefs of the same subject.¹
- (2) To say that a belief is immediately justified is to say that what justifies it does not include some other justified beliefs of the same subject.

This generic characterization of immediate justification is purely negative. Anyone who holds that some beliefs are immediately justified will have some conception of what can justify beliefs in such a way that no other justified beliefs of the same subject are involved in the justification.

Now the confusion about immediate justification I will be exploring consists just in this: it is confusedly supposed that for S's belief that *p* to be immediately justified it is required that the higher-level belief that *S is justified in believing that p*, or that *S knows that p*, itself be immediately justified; or, even more confusedly, that this is what the immediate justification of S's belief that *p* consists in. Full-blown examples of this confusion can be found in Roderick Chisholm and in Panayot Butchvarov.² I will restrict my attention to Chisholm.

Chisholm's version of immediate justification is what we may call truth-justification, justification of a belief by its truth or by the fact that makes it true. To follow Chisholm's presentation of this, a short terminological digression will be required. Chisholm distinguishes several grades of epistemic justification, one of the higher of which is 'evident'. (The exact definition of 'evident' and its distinction from other grades need not concern us here.) The term 'evident' is applied to propositions; if a proposition, *p*, is evident for a subject, S, then S is justified (to a high

degree) in believing that p . Chisholm tends to use the term 'justified' in a non-discriminating way to range over all grades of justification.

In the recently published second edition of his *Theory of Knowledge*³ Chisholm defines his basic notion of immediate justification for empirical beliefs as follows.

D2.1 b is *self-presenting* for S at $t = Df$. b is true at t ; and necessarily, if b is true at t , then b is *evident* for S at t . (p. 22)⁴

This conforms to the generic notion of immediate justification I presented above. But Chisholm also presents his version of immediate justification in a quite different way. He introduces his conception of the directly evident by considering the ways in which one might answer the "Socratic" questions "What justification do you have for thinking you know this thing to be true?" or "What justification do you have for counting this thing as something that is evident?" (p. 17)

In many instances the answers to our questions will take the following form: "What justifies me in thinking that I know that a is F is the fact it is evident to me that b is G " . . . This type of answer to our Socratic questions shifts the burden of justification from one claim to another. For we may now ask, "What justifies me in counting it as evident that b is G ?" or "What justifies me in thinking I know that b is G ?" . . . We might try to continue *ad infinitum*, justifying each new claim that we elicit by still another claim. Or we might be tempted to complete a vicious circle . . . But if we are rational beings, we will do neither of these things. For we will find that our Socratic questions lead us to a proper stopping place . . . Let us say provisionally that we have found a proper stopping place when the answer to our question may take the following form:

What justifies me in thinking I know that a is F is simply the fact that a is F .

Whenever this type of answer is appropriate, we have encountered the *directly evident*. (pp. 18-20)

In this passage and others we get a different picture of what makes a proposition directly evident. According to the definition D2.1, what makes a true proposition, p , directly evident for S , is that its truth makes it evident for S ; whereas according to the passage just quoted what makes p directly evident is that its truth makes evident (justifies)⁵ S 's higher level belief that S knows that p (or that it is evident to S that p). The two passages give different answers to the question: what does the truth of p have to justify in order that p be *directly* evident?

There is fairly strong textual evidence that Chisholm simply does not see that the two accounts are different, or, at least, that the realization of their difference is not effectively operative in his mind when he is presenting his position. Not only do we find each account reflected in numerous passages. We even find Chisholm juxtaposing them in the same discussion.

Thinking and believing provide us with paradigm cases of the directly evident.

Consider a reasonable man who . . . believes that Albuquerque is in New Mexico, and suppose him to reflect on the philosophical question, "What is my justification for thinking that I know . . . that I believe that Albuquerque is in New Mexico? . . . The man could reply in this way: "My justification for thinking I know . . . that I believe that Albuquerque is in New Mexico, is simply the fact . . . that I do believe that it is in New Mexico." And this reply fits our formula for the directly evident:

What justifies me in thinking I know that a is F is simply the fact that a is F .

Our man has stated his justification for a proposition merely by reiterating that proposition. (p. 21)

Obviously it is the higher-level conception of direct evidence that is being employed throughout most of this passage. But the very last sentence constitutes a reversion to the lower-level conception. If the proposition for which the man is stating his justification was the higher-level proposition *I know that I believe that Albuquerque is in New Mexico*, then he did not state his justification by reiterating the proposition. For what he enunciated in stating his justification was not that proposition, but its lower-level correlate, *I believe that Albuquerque is in New Mexico*. Thus he stated his justification for p by reiterating p only if the p in question were that lower-level proposition.

Of course it may be that Chisholm is not confusing the two levels but is presenting the matter in such a way as to reflect his *conviction* that, for self-presenting propositions, the truth of p generates justification on both levels. Indeed, in a later part of the book Chisholm does espouse, and argue for, a level-bridging principle that might seem to have this consequence.

. . . if a proposition is evident and if one considers the proposition, then it is evident that the proposition is evident. (p. 114)

This principle does ensure a transfer of evidence from a proposition, p , to the higher-level proposition that *it is evident that p*, given that S considers the matter. But it by no means follows from this that the source of evidence is the same on the two levels; hence it does not follow that where the truth of p suffices to make p evident, *it will also suffice to make it is evident that p evident*. The principle quoted above is quite compatible with its being the case that where *it is evident that p* (for some self-presenting proposition, p) becomes evident to S upon considering the matter, what makes the higher-level proposition evident is not the mere truth of p , but something that is uncovered in the process of consideration. And Chisholm evinces no awareness that the thesis that the truth of p generates evidence on the higher level as well as the lower, is one that needs to be scrutinized and defended, whether on the basis of the above principle or otherwise.

In any event, the important philosophical question is not what is or is not going on in Chisholm's mind, but whether the thesis that the source of evidence

is the same on the two levels, has important consequences that are likely to pass unnoticed if one simply assumes the thesis without explicitly realizing that one is doing so. I will now point out some of those consequences.

First, if one saddles one's account of immediate justification with the claim that the same kind of justification extends to one or more correlated higher-level propositions, the plausibility of one's account will be reduced. This is certainly the case with Chisholm. Whatever ultimate judgment is to be brought in the matter, it is not totally implausible to suppose that one is justified in beliefs about what one is currently feeling, sensing, or thinking just by the fact that one is so feeling, sensing, or thinking. But is it equally plausible that I am justified in supposing that *it is evident to me that I feel tired* just by the fact that I feel tired? Can I be justified in supposing that a certain proposition has a certain epistemic status for me, *just by feeling tired*. One's initial doubts in this matter are increased by considering Chisholm's definition of 'evident'.

D1.5 *b* is evident for *S* = Df (i) *b* is beyond reasonable doubt for *S* and (ii) for every *i*, if accepting *i* is more reasonable for *S* than accepting *b*, then *i* is certain for *S*. (p. 12)

And the definition of 'certain' runs:

D1.4 *b* is certain for *S* = Df *b* is beyond reasonable doubt for *S*, and there is no *i* such that accepting *i* is more reasonable for *S* than accepting *b*. (p. 10)

Leaving aside what it takes to be justified in supposing the acceptance of one proposition to be more reasonable than the acceptance of another, and leaving aside what it takes to be justified in supposing that a certain proposition is beyond reasonable doubt, let us concentrate on the rest of what is involved in a proposition's being evident, viz., a certain comparative epistemic status vis-à-vis all other propositions. More specifically, this comparative status consists in its being the case that no other propositions enjoy a more favorable epistemic status for *S* except those that enjoy the highest possible epistemic status. Now, is it credible that I should be justified in a belief that is, in part, about the epistemic status of a given proposition vis-à-vis the entire class of propositions, *just by virtue of feeling tired*? At the very least, the claim to higher-level truth-justification raises questions that are quite different from the claim to lower-level truth-justification. Chisholm has saddled his theory with a considerable liability by adding on the higher-level claim.⁶

Moreover, Chisholm need not have taken on this additional liability in order for direct evidence to play its intended role in his system. The course of Chisholm's exposition, and the structure of his theory, makes it clear that the main function of directly evident propositions in his system is to stop the regress of justification and serve as foundations of knowledge. I have argued elsewhere that the demands of the regress argument are amply satisfied by first-level immediate justification and that a foundationalist epistemology based on propositions that enjoy only first-level immediate justification will be in at least as strong a position as any other foundationalism.⁷ It is true that Chisholm's methodology requires what we might call

"high accessibility" to one's own epistemic states. This position is reflected in the quote from p. 114 given above and in other pronouncements in that same section of the book, such as Chisholm's version of the KK thesis.

(K4) If *S* considers the proposition that he knows that *p*, and if he does know that *p*, then he knows that he knows that *p*. (p. 116)

However, it remains to be shown that high accessibility requires that what justifies the higher-level proposition that *it is evident to S that p*, or that *S knows that p*, be the *same* as what justifies *p* itself. Chisholm has not so argued, and I am dubious about the prospects.

An equally serious consequence of a confusion of levels (or of an uncritical assumption that correlated propositions on two levels enjoy the same justification) is that the range of candidates for immediate justification is sharply restricted. It is a striking fact that most epistemologists who work with something like our distinction between mediate and immediate justification are markedly penurious in the modes of immediate justification they consider. Chisholm is typical in this regard. He simply notes that when a proposition is rendered evident by its own truth it is thereby *directly* evident, and he fails to consider whether there are other possibilities. Other epistemologists are equally narrowly preoccupied with immediate awareness or with self-evidence as sources of immediate justification.⁸ One particularly unfortunate consequence of this parochialism is an obliviousness to the possibility that a belief might be immediately justified by having originated in a certain way, e.g., justified by having been produced by a reliable belief-producing mechanism.⁹ Whatever the reason for Chisholm's ignoring immediate awareness, or Lewis's ignoring truth-justification, it seems quite plausible to suppose that the level-confusion we have been discussing is responsible for the widespread neglect of immediate justification by origin. For if one takes it that *S* is immediately justified in believing that *p* only if *S* is immediately justified in believing that *S is justified in believing that p*, then one will restrict the range of immediate justifiers to those one supposes will be capable of justifying the higher-level, as well as the lower-level, proposition. As we have seen in discussing Chisholm, it is by no means obvious that the modes of immediate justification favored by level-confusers do meet this requirement; perhaps a judicious assessment would reveal that none do. Nevertheless, it seems much *more* obvious that the fact that a belief was produced by a reliable psychological mechanism is *not* sufficient to justify a belief *about* the epistemic status of that belief; for we are often in the dark concerning the reliability, or other features, of what produces our beliefs. Hence in failing to distinguish between justification on the two levels, one will be led to ignore the possible epistemic relevance of the actual mode of belief-production.

Indeed, even where the possibility is considered, level-confusions may play a decisive role in its evaluation. Consider the following passage from Keith Lehrer's book, *Knowledge*.¹⁰

Thus, if something looks red to a person, he cannot justifiably conclude that it is red from the formula that red things look red in standard conditions to

normal observers, he would also need to know that the conditions are standard and that he is normal. Independent information is, therefore, required for the justification of this perceptual belief. . . . More generally, to justify such a belief requires the information that the conditions that surround a man and the state he is in are such that when something looks red in conditions of this sort to a person in his state, then it is red.

. . . Since a man may hallucinate, he cannot justifiably conclude he sees something as opposed to merely hallucinating unless he has information enabling him to distinguish hallucination from the real thing. (pp. 103-4)

Let us agree that a person to whom x looks red cannot be justified in a perceptual belief that x is red unless "the conditions that surround" him and "the state he is in are such that when something looks red in conditions of this sort to a person in his state, then it is red." But why should we also require that the person *have that information, know* (justifiably believe) that this is so. Why is it not enough that it be so? As we read on, it becomes transparently clear that Lehrer is falling into a level-confusion.

. . . the need for independent information arose from the need to determine whether the circumstances in which a person finds himself are those in which a man may justifiably conclude that he is seeing a typewriter or seeing something red.

. . . when a great deal . . . hinges on the matter of whether the person saw a bear-print or something else, . . . then we start to ask serious questions. We seek to determine if the person has information enabling him to decide whether he is seeing things of the sort he sees. (p. 105)

Well of course if *that* is what we are (he is) after, we (he) need "independent information." If he is trying to determine whether he is (really) seeing a bear-print (which involves determining whether his perceptual belief that there was a bear-print in a certain place was justified), or trying to determine whether the circumstances of his perception were such as to justify his perceptual belief, then *of course* he needs evidence of the sort mentioned. But that is just to say that he needs such evidence in order to be justified in the higher-level epistemic belief that his original perceptual belief was justified (and to be justified in the beliefs that support that epistemic belief). Lehrer can get from this incontrovertible truth to his central claim that such information is required for the perceptual belief to be justified, only by confusing the two problems — the justification of the perceptual belief and the justification of the higher-level belief that the perceptual belief is justified.

If one restricts oneself to sources of immediate justification that one supposes, survive a transition to higher levels, the kinds of beliefs one takes to be susceptible of immediate justification will be likewise restricted. Historically, this has meant a restriction (for *a posteriori* knowledge) to beliefs concerning the believer's current states of consciousness. The insuperable difficulties encountered in the attempt to build the whole of *a posteriori* knowledge on such a slim basis have been more than amply documented. Our discussion reveals the role level-confusion

has played in generating the supposition that no more extended foundation is available.

Indeed, if one does not distinguish between justification on different levels, one may be, confusedly, led to reject the whole concept of immediate justification. Consider the following argument from Bruce Aune's book, *Knowledge, Mind, and Nature*.¹¹

I would venture to say that any spontaneous claim, observational or introspective, carries almost no presumption of truth, when considered entirely by itself. If we accept such a claim as true, it is only because of our confidence that a complex body of background assumptions — concerning observers, standing conditions, the kind of object in question — and, often, a complex mass of further observations all point to the conclusion that it is true.

Given these prosaic considerations, it is not necessary to cite experimental evidence illustrating the delusions easily brought about by, for example, hypnosis to see that no spontaneous claim is acceptable wholly on its own merits. On the contrary, common experience is entirely adequate to show that clear-headed men never accept a claim merely because it is made, without regard to the peculiarities of the agent and of the conditions under which it is produced. For such men, the acceptability of every claim is always determined by inference. If we are prepared to take these standards of acceptability seriously, we must accordingly admit that the traditional search for intrinsically acceptable empirical premises is completely misguided. (pp. 42-43)

Here Aune is arguing that beliefs are justified only by inference (from other propositions known, or justifiably believed), which is equivalent to the denial that there are any immediately justified beliefs. But a close reading will reveal that the considerations he advances seem to yield that conclusion only if one is confusing levels. The solid points that Aune makes in support of that claim are the following.

If we accept such a claim [observational or introspective] as true, it is only because of our confidence that a complex of background assumptions . . . all point to the conclusion that it is true.

. . . clear-headed men never accept a claim merely because it is made, without regard to the peculiarities of the agent and of the conditions under which it is produced. For such men, the acceptability of every claim is always determined by inference.

Now in making these points Aune is not really considering what would justify the *issuer* of an introspective or observational claim, but what it would take to justify "us" in accepting his claim; he is considering the matter from a third-person perspective. And it is clear that I cannot be immediately justified in accepting *your* introspective or observational claim. If I am so justified, it is because I am justified in supposing that you issued a claim of that sort, that you are in a normal condition and know the language, and so on. But that is only because I, in contrast to you, am justified in believing that *p* (where what you claimed is that *p*) only if I am

justified in supposing that *you are justified in believing that p*. My access to *p* is through your access. It is just because my justification in believing *p* presupposes my being justified in believing that you are justified, that my justification has to be indirect. Thus what Aune's argument supports is the necessity for inferential backing for any higher-level belief to the effect that some person is justified in believing that *p*. Only a failure to distinguish levels leads him to suppose that he has shown that *no* belief can be immediately justified.

II

Next let us consider the bearing of level-confusions on the requirements for *mediate* justification, or in less technical terminology, on what it takes for *S* to have an adequate reason, grounds, or evidence for supposing that *p*. If the justification is *mediate*, there must be some other proposition, *q*, that is related to *p* and to *S*'s belief that *p* in certain ways. Exactly what ways are necessary? The following requirements are accepted by virtually all who have considered the matter.

- (1) *q* is related to *p* in a way that is "appropriate"¹² for purposes of justification.
- (2) *S* believes that *q*.
- (3) *S* is justified in believing that *q*.¹³

Most of the discussion of *mediate* justification has centered around (1). How must propositions, e.g., about sensory appearances, be related to, e.g., propositions about physical objects in the environment of the perceiver, to serve as adequate grounds for the latter? Must there be an entailment? Will some sort of inductive evidence relationship do? Or is there some special "evidence-conferring" relationship involved?

Again, there is widespread agreement that there must be some "psychological" connection between *S*'s belief that *q* and *S*'s belief that *p*. They cannot just lie "side by side" in his mind: *q* must be "his reason," or at least one of his reasons for believing that *p*. This is often taken to imply that the belief that *p* have been produced by the belief that *q*, or that the former be causally *sustained* by the latter. Sometimes this is further specified to require that *S* have *inferred* *p* from *q*, or now be disposed to do so. But whether or not inference is required, there is general agreement that some restrictions must be put on the mode of generation. So let us put as the fourth condition:

- (4) *S*'s belief that *p* was produced by, or is causally sustained by, *S*'s belief that *q*, in the right way.

Now we come to further alleged conditions that, I want to suggest, depend for their plausibility on level-confusions. For one thing, various writers¹⁴ hold that if *S*'s belief that *q* is to constitute an adequate basis for *S*'s belief that *p*, not only must *q* be appropriately related to *p*, but *S* must *know*, or at least justifiably believe, that this is so.

(5) S is justified in believing that *q* is appropriately related to *p*.

It seems to me that this is too sophisticated as a general requirement for mediate justification, especially if we take mediate justification to be required for mediate knowledge. Surely creatures like dogs and preverbal children can have mediate knowledge. My dog knows that I am preparing to take him for a walk, and he knows that because he sees me getting out his chain. But such creatures have no concepts of deductive, inductive, or other relations between propositions, and hence are quite incapable of believing, much less justifiably believing, that such relations obtain. Even where S has the relevant concepts, he may not be *justified* in supposing that appropriate relations obtain. He may just unthinkingly assume (truly) that, e.g., his local newspaper is a reliable source of local news. Does this prevent him from learning (coming to know) about local happenings from reading his newspaper (from his knowledge that these happenings are reported in the newspaper)?

Those who introduce condition (5) fail to give anything like a full-dress defense of it. Its proponents seem to take it as having sufficient intrinsic plausibility to make an explicit defense unnecessary. My diagnosis is that this plausibility largely stems from level-confusion. It does seem that I cannot be justified in the higher-level belief that *my belief that q mediately justifies me in believing that p* unless I am justified in supposing that *q* is appropriately related to *p*. For unless I am justified in supposing that, how could I be justified in supposing that the appropriate justification relation holds between the beliefs? And so if one does not distinguish between being justified in believing that *p* and being justified in *supposing that one is mediately justified in believing that p*, then one will naturally suppose that what is required for the latter is also required for the former.

Another widespread requirement is:

(6) S is able, or disposed, to cite *q* as what justifies his belief that *p*.

Here, e.g., is C. I. Lewis, disavowing the necessity for a conscious inference from *q* to *p*, and replacing that requirement with a combination of (4) and (6).

. . . whether the ground of judgment is or is not explicitly in mind, is hardly the pertinent consideration, because it could not plausibly be taken to mark the important distinction between attitudes of B having positive cognitive value and those which lack it. Rather the pertinent distinction is between cases in which if the judgment be challenged by ourselves or others, we should be able to assign a basis of it which, whether explicitly thought of in drawing the judgment or not, is so related to it that we could truly say "If it were not for that, I should not have so judged."¹⁵

Again (6) would seem to be much too sophisticated a requirement, especially if justification is required for knowledge. There are knowing creatures who lack the sophistication, or even the linguistic skills, to respond to challenges by specifying the basis of their beliefs. They include creatures that do not have the use of language as well as language users who do not (yet) have any concept of epistemic

justification. Even those sophisticated enough to engage in this kind of palaver may be unable, in particular cases, to identify the real and sufficient bases of their belief. Why, then, has this requirement seemed right to many? Here too level-confusion may play an important role. Requirement (6) seems more plausible as a requirement for being justified in accepting the higher-level proposition that *S is justified in believing that p*. One might well think that I cannot be justified in a claim to justification unless I can point out what does the justifying. But here we cannot pin all the blame on level-confusion. For, in truth, (6) is questionable as a requirement for higher-level justification as well. Why must I be able to *specify, cite, or formulate* what it is that justifies me in believing that *p*, in order to be justified in supposing that I *am* so justified? This is a special form of the old question of whether I can be justified in accepting a relatively unspecific or general proposition without being able to specify the particular fact(s) that makes it true. It has many forms: can I not be justified in supposing that there are a lot of dots on that surface without being able to say how many? Can I not be justified in believing that there is someone in the room without being able to say who is in the room? Of course it remains to be seen exactly how one could be justified in supposing, unspecifically, that he is (somehow) justified in believing that *p* without being able to say precisely what justifies him. But surely this possibility should not be dismissed without a hearing.

In the light of the point just made, perhaps the main villain in this piece is another widespread confusion in epistemology — one we are not really exploring in this paper — the confusion between ‘justification’ in the sense of *being* justified and ‘justification’ in the sense of “showing that one is justified.” If one fails to keep that distinction in mind, one is liable to suppose that in order to *be* justified in believing that *p* one must be *able* at least to “justify” one’s belief that *p* in the sense of showing that one is justified, i.e., exhibiting what it is that justifies one. And that would explain the plausibility of (6).

III

Finally let us consider the role of level-confusion in certain forms of skeptical argument. First, look at what may conveniently be called “Cartesian skepticism” because of its similarity to what we find in Descartes’s *Meditations*.¹⁶ The kind of argument I wish to discuss is directed at some particular knowledge claim and is designed to show that the claimer, *S*, does not know what he is claiming to know. Let us consider a case in which a person is looking out the window and claims to know that a car is parked in front of his house. (He supposes himself to see a car parked there.) The argument will then proceed as follows.

1. If *S*’s present visual experience is being directly produced by an omnipotent spirit, then *S* does not know (perceptually) that there is a car parked in front of his house.¹⁷
2. *S* does not know that his present visual experience is not being directly produced by an omnipotent spirit.

3. Therefore, S does not know (perceptually) that there is a car parked in front of his house.

Questions could be raised about both premises, but I will not go into that. Instead, I will contend that even if both premises were unexceptionable, the conclusion would not follow. Why should we suppose that S's inability to rule out the hypothesis of an abnormal production of his visual experience implies that his visual experience gives him no knowledge of the physical environment? Any answer to this question will have to derive from our rationale for (1). Let us take that rationale to depend on some kind of (at least partly) causal theory of perceptual knowledge. My visual experiences can give me knowledge of a certain physical object only if that object played a role in the chain of causes leading up to that experience. If those experiences would have been produced exactly as they were (given the particular circumstances in which they occurred) even if that object were not there, then those experiences cannot mediate any knowledge of that object. If this be accepted, then (1) is justified. If S's visual experiences were produced directly by an omnipotent spirit, then they would have been produced in precisely this form even if a car had not been parked in front of his house. Hence, in that case, he would not know in this way, would not have visual knowledge, that there is a car parked in front of his house. But how does the conclusion follow from all that (plus (2)). Granted that an *actual* abnormal production inhibits perceptual knowledge, why suppose that the mere fact that S does not know the production was not abnormal rules out S's knowing about the car? If the object I am eating is made of cardboard, it will not nourish me. But suppose I do not know it is not made of cardboard; it by no means follows just from this lack of knowledge that the object will not nourish me. Its nutrient power, or the reverse, depends on what it *is*, not on what I do or do not *know* about it. Why should we suppose the present case is any different?

Here is a slightly different way of putting the matter. I do not know whether what I am eating is made of cardboard. But that fact leaves wide open the possibility that it is not made of cardboard and that it in fact contains nutrients. Similarly, the fact that I do not know whether my present visual experiences are being directly produced by an ingenious neuro-physiologist leaves wide open the possibility that in fact they are being produced in the usual way by a chain of causes stemming from a car parked in front of my house. And if that possibility is realized, I do have perceptual knowledge that a car is parked in front of my house. Since premise (2) does not rule out the possibility in question, it (with premise (1)) does not establish that I do not know that a car is parked in front of my house.

But then why is this argument so tempting? Again, a level-confusion may be largely responsible. Given a certain assumption, we can derive a higher-level correlate of (3) from our two premises, a correlate that replaces *there is a car parked in front of S's house* with *S knows (perceptually) that there is a car parked in front of S's house*.

3A. Therefore S does not know that he knows (perceptually) that there is a car parked in front of his house.

The assumption in question, a rather controversial one, is that one cannot know that p unless one knows, with respect to each of the necessary conditions of p , that it obtains. Now according to premise (1), one necessary condition of S's knowing (perceptually) that there is a car parked in front of his house is that his perceptual experience is not produced abnormally. But according to (2), S does not know that this necessary condition obtains. Hence (3A): he does not know that he knows (perceptually) that there is a car parked in front of his house. But, granted that (3A) follows from (1) and (2), why suppose that (3) follows? One possible explanation of this supposition is a *conviction* that one cannot know that p without knowing that one knows that p ; if that were so, then to show that one does not know that one knows that p is *ipso facto* to show that one does not know that p . However, not many philosophers hold so strong a level-bridging view. Hence I think that the attractiveness of the original argument is largely due to a level-*confusion*. If one fails to distinguish clearly between p and S knows that p , one will likewise not distinguish between what it takes to know the one and what it takes to know the other.¹⁸

Finally, let us consider another kind of skeptical argument, in which level-confusion also plays an important part. This is what we may call "criterion skepticism"; the classical form is in Sextus Empiricus, *Outlines of Pyrrhonism*, bk. II, chap. 4.

In order to decide the dispute which has arisen about the criterion (of truth), we must possess an accepted criterion by which we shall be able to judge the dispute; and in order to possess an accepted criterion, the dispute about the criterion must first be decided. And when the argument thus reduces itself to a form of circular reasoning the discovery of the criterion becomes impracticable, since we do not allow them to adopt a criterion by assumption, while if they offer to judge the criterion by a criterion we force them to a regress *ad infinitum*. And furthermore, since demonstration requires an approved demonstration, they are forced into circular reasoning.

I should like to work with my own version of an argument suggested by these remarks of Sextus.

In order for me to be justified in believing that p , my belief that p must satisfy the conditions laid down by some valid epistemic principle (for epistemic justification). But then I am justified in the original belief only if I am justified in supposing that there is a valid epistemic principle that does apply in that way to my present belief. And in order to be justified in that further belief there must be a valid epistemic principle that is satisfied in *that case*. And in order to be justified in supposing that . . . This series either doubles back on itself, in which case the justification is circular, or it stretches back infinitely. Thus it would appear that claims to justification give rise either to circularity or to an infinite regress.

The level-confusion is more readily apparent here than in Cartesian skepticism. This argument has no tendency to show that my being justified in believing that p

depends on conditions that give rise to an infinite regress. On the argument's own showing, what my *being* justified in believing that *p* depends on is the existence of a valid epistemic principle that applies to my belief that *p*. So long as there is such a principle, that belief *is* justified whether I know anything about the principle or not and whether or not I am justified in supposing that there is such a principle. What this latter justification is required for is not my being justified in believing that *p*, but rather my being justified in the higher-level belief that *I am justified in believing that p*. I can be justified in that higher-level belief only if I am justified in supposing there to be a principle of the right sort. But it is only by a level-confusion that one could suppose this latter justification to be required for my being justified in the original lower-level belief. The regress never gets started.

This would seem to leave open the possibility that being justified in a higher-level belief, such as the belief that *I am justified in believing that p*, does give rise to an infinite regress or circularity. But that would be a mistake of the same kind. To be justified in that higher-level belief, there has to be a (higher-level) epistemic principle of justification that applies in the right way to the belief in question. But again, all that is required is the *existence* of such a principle. For the justification of that (first-order) higher-level belief, it is not necessary that I be justified in supposing that there is such a principle; only that there be such. Again, what this last justification is needed for is the justification of the still higher-level belief that *I am justified in believing that I am justified in believing that p*. At each stage I can be justified in holding a certain belief provided there is a valid epistemic principle that satisfies certain conditions. My knowing or being justified in believing that there is such a principle is required only for the justification of a belief that is of a still higher level vis-à-vis the belief with which we started.

IV

In conclusion, let me suggest a more positive moral from this string of polemics. It should be clear that the level-confusions we have been examining naturally lead to ignoring the possibility of what we might call unsophisticated, unreflective first-level knowledge or justification, cases in which one knows that *p*, or is justified in believing that *p*, but, whether because of conceptual underdevelopment or otherwise, fails to attain the more sophisticated, higher-level knowledge (or justified belief) that one has that lower-level knowledge or justification. Of course it may not be immediately obvious that there is unreflective knowledge or justification; the question needs careful consideration. But the point is that so long as we are victims of level-confusion we cannot even consider the possibility of a purely first-level cognition. The new look in epistemology introduced by the "reliability" theories of such thinkers as Dretske, Armstrong, and Goldman is largely built on the claim that first-level knowledge is independent of higher-level knowledge. We will be able to take this "new look" even experimentally, only to the extent that we can free ourselves from the blinders imposed by level-confusion.

Notes

1. We are leaving open the question of what else is required for mediate justification over and above the possession of certain other justified beliefs.
2. Panayot Butchvarov, *The Concept of Knowledge* (Evanston, Ill., 1970), pt. I, section 6.
3. Roderick Chisholm, *Theory of Knowledge*, 2nd ed. (Englewood Cliffs, N.J., 1977).
4. One other terminological guide to the quotations that follow. In the first edition of *Theory of Knowledge* (1966), Chisholm used the term 'directly evident' for the concept expressed above by the term 'self-presenting'. In the second edition the former term is officially reserved for a wider concept, but much of the material that was retained from the first edition still used 'directly evident' in the way 'self-presenting' was defined above. I will use the term 'directly evident' for the concept just defined.
5. What are we to make of the fact that in D2.1 Chisholm speaks of *p* being made *evident* by the fact that *p*, whereas in the passage just quoted he speaks of the higher-level proposition being *justified* by the fact that *p*? Does this indicate that Chisholm is less sure about the degree of justification conferred by the fact that *p* on the higher-level proposition than on the degree of justification it confers on the proposition that *p*? Or does he think that the propositions on both levels are made *evident* by the fact that *p*? For present purposes it is not necessary to settle this question. The point with which we are concerned is simply the relation between claims to some justificatory role of the fact that *p* on the two levels.
6. Similar points can be made for other conceptions of immediate justification. If we hold with Russell, C. I. Lewis, and many others, that beliefs about one's current sensory data are justified by the fact that one is "directly aware" of those states, this has a certain initial plausibility, one that is not shared by the correlated higher-level claim that one is justified in believing that one is justified in holding such beliefs by virtue of being directly aware of sensory data.
7. William P. Alston, "Two Types of Foundationalism", *Journal of Philosophy* 73 (1976): 165-85.
8. For two of the rare attempts to critically compare different putative direct justifiers, see my essay "Self-Warrant: A Neglected Form of Privileged Access," *American Philosophical Quarterly* 13 (1976): 257-72; Butchvarov, *The Concept of Knowledge*, Chap. 1, sec. 6.
9. For some presentations of this possibility see Alvin I. Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy* 73 (1976): 771-91; and my "The Justification of Perceptual Beliefs," unpublished.
10. Keith Lehrer, *Knowledge* (Oxford, 1974).
11. Bruce Aune, *Knowledge, Mind, and Nature* (New York, 1967).
12. If one should try to give a general criterion for "appropriateness," it might be something like this: *q* is related to *p* "appropriately" iff the truth of *q* will thereby either guarantee the truth of *p*, or at least make the truth of *p* likely. In other words, the relationship is, or tends to be, truth-preserving.
13. The rationale for (2) and (3) are fairly obvious. How can the fact that *q* is "appropriately" related to *p* do anything to justify *me* in believing that *p* unless I "have" this adequate ground, unless I am in a position to appropriate the epistemic benefits contained therein. And I cannot do this unless it is at least something I believe. And unless I am *justified* in believing it, how can justification (for me) be transferred along the appropriate propositional relation. Some would go further and require that I *know* that *q*. The temporal relations between the beliefs are widely ignored.
14. D. M. Armstrong, *Belief, Truth, and Knowledge* (Cambridge, 1973), p. 151. Brian Skyrms, "The Explication of 'X knows that *p*'" *Journal of Philosophy* 64 (June 22, 1967): 374.
15. C. I. Lewis, *An Analysis of Knowledge and Valuation* (LaSalle, Ill., 1946), p. 328.
16. This argument is not supposed to be an exact replica of anything in the *Meditations*.

17. For a more up-to-date version the omnipotent spirit could be replaced by an ingenious neuro-physiologist.

18. Of course, if (1) were of the form 'If q , then not- p ', rather than of the form 'If q , then S doesn't know that p ', it would be a different ball game. (Where p is, in our case, *There is a car parked in front of S's house*, and q is S 's present visual experience is being directly produced by an omnipotent spirit.) For in that case the falsity of q is one of the necessary conditions of the truth of p , and so (2) tells us that S does not know that this necessary condition holds. And so the same reasoning that led us to take the original argument to show that S does not know that S knows that p , would lead us to take this argument to show that S does not know that p . Sometimes Cartesian skepticism is presented in this stronger form and sometimes in the weaker form. Thus when q is *I am dreaming* and p is *I am seated in front of the fire awake*, we have the stronger form, for q does imply not- p . But in our original example, q did not imply not- p . My present visual experience's being produced by an omnipotent spirit is quite compatible with there being a car parked in front of my house at the moment. In this paper I am concerned only with the weaker form. It is worthy of note that the stronger form is more vulnerable to the Moore-Malcolm charge of begging the question. For if q does imply not- p , then the question of whether I know that not- q is directly dependent on whether I know that p . For if I do know that p , which is the point of contention, then, given certain principles of epistemic logic, I *ipso facto* know that not- q .