

3

INTERNALISM AND EXTERNALISM IN EPISTEMOLOGY

William P. Alston

One hears much these days of an epistemological distinction between "internal" and "external." It is often found in discussions of reliabilism in which the critic accuses the reliabilist of violating "internalist" restrictions on justification and of resting content with justification that is "external" to the subject's perspective.¹ But just what distinction is this (are these)? That is not so clear.

As just intimated, those who wield the distinction intend to be contrasting different views on what can confer justification or on what can convert mere true belief into knowledge. The main emphasis has been on justification, and we will continue that emphasis in this paper. In all these discussions it is the internalist position that lays down constraints; the externalist position *vis-à-vis* a given internalist position is simply the denial that the internalist constraint in question constitutes a necessary condition of justification. Thus our attempts at clarification can be confined to the internalist side.

As the name implies, an "internalist" position will restrict justifiers to items that are *within* something, more specifically, within the subject. But, of course, not everything that is "within" a knowing subject will be admitted as a possible justifier by an internalist. Physiological processes within the subject, of which the subject knows nothing, will not be allowed. Then just where, how, or in what sense, does something have to be "in the subject" in order to pass the internalist test?

Two quite different answers are given to this question in the literature. First there is the idea that in order to confer justification something must be within the subject's "perspective" or "viewpoint" on the world, in the

sense of being something that the subject knows, believes, or justifiably believes. It must be something that falls within the subject's ken, something of which the subject has taken note. Second there is the idea that in order to confer justification something must be accessible to the subject in some special way, e.g., directly accessible or infallibly inaccessible. We shall explore each of these versions in detail, noting alternative formulations of each, exposing unclarity and incoherences, and seeking to develop the strongest form of each position. We shall consider what can be said for and against each version, and we shall explore their interrelations. Finally we shall make some suggestions concerning the most reasonable position to take on these issues.

I

Let's begin by considering the first form of internalism. In the essay already cited, BonJour, in discussing the view that there are "basic beliefs", has this to say:

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

The specific conclusion here is that there can be no basic beliefs, no beliefs that are justified otherwise than on the basis of other beliefs. But that is not our present concern. We are interested in the constraint on justification invoked by BonJour to arrive at this result. That is the requirement that "that feature, whatever it may be, in virtue of which an empirical belief qualifies as basic", that is, that feature by virtue of which it is justified, must be justifiably believed by the subject to attach to that belief if the belief is to be thereby justified. That is, the justifying feature must be part of his "perspective on the world", must be known or justifiably believed by him to obtain if it is to do its justifying work.

BonJour continues to employ this same understanding of internalism in characterising the opposed externalist position.

But according to proponents of the view under discussion, the person for whom the belief is basic need not (and in general will not) have any cognitive grasp of any kind of this reason or of the relation that is the basis for it in order for this basic belief to be justified; all these matters may be entirely *external* to the person's subjective conception of the situation.³

When viewed from the general standpoint of the western epistemological tradition, externalism represents a very radical departure. It seems safe to say that until very recent times, no serious philosopher of knowledge would have dreamed of suggesting that a person's beliefs might be epistemically justified simply in virtue of facts or relations that were external to his subjective conception.⁴

Again, in "A Rationale for Reliabilism" Kent Bach writes as follows.

Internalism requires that a person have "cognitive grasp" of whatever makes his belief justified.⁵

And in "The Internalist Conception of Justification" Alvin Goldman writes:

Traditional epistemology has not adopted this externalist perspective. It has been predominantly *internalist*, or egocentric. On the latter perspective, epistemology's job is to construct a doxastic principle or procedure *from the inside*, from our own individual vantage point.⁶

All this would suggest the following formulation of internalism.

- (1) Only what is within the subject's "perspective" can determine the justification of a belief.

Let's call this version of internalism "perspectival internalism" (henceforth "PI").

PI needs some refinement before we are ready to consider what can be said for and against it. First, we have been specifying the subject's "perspective" disjunctively as what the subject "knows, believes, or justifiably believes". It will make a considerable difference what choice we make from between these alternatives. For the present let's proceed in terms of justified belief. At a later

stage of the discussion we will explicitly consider the three alternatives and justify this decision. This gives us the more specific formulation:

- (2) Only the justified beliefs of the subject can determine what further beliefs of that subject are justified.

(2) may seem to smell of circularity, but there can be no definitional circularity, since the internalism we are discussing is not concerned with defining "justified"; it is merely laying down one constraint on the provision of justification. There are, of course, well known problems with making all justification depend on other justified beliefs, and we shall attend to these in due course.

Next we need to consider the way in which the perspective *determines* the justification of belief. But first a terminological matter. BonJour's formulation is in terms of a "feature" of the belief by which it is justified. Sometimes this is the most natural construal, as when we think of beliefs about ones current conscious states as being justified by virtue of the fact that they, the beliefs, are incorrigible, or by virtue of the fact that they, the beliefs, are "self-warranted". However it is usually more natural to think of the justification of a belief as stemming from its relation to some state of affairs other than itself, as when a belief is justified by virtue of being based on adequate evidence or reasons, or by virtue of arising from a certain sensory experience. To be sure, these ways of talking are mutually translatable. By a well known grammatical trick we can always take a belief's relation to some external justifying state of affairs to be a property of the belief. And, contrariwise, we can take the fact that belief B is incorrigible to be the state of affairs that justifies it. Hence I shall feel free to use now one construal, now the other, as seems most natural in the particular context. However I will most often speak in terms of justifying *facts* or *states of affairs* and will refer to them as "justifiers".

Let's return to the issue concerning the way in which the perspective determines justification. In the first quotation from BonJour he allows any sort of fact, not just other justified beliefs of the subject, to be a justifier, provided the subject has certain justified beliefs concerning it and its relation to the initial belief. A justifier for a perceptual belief that there is a tree in front of one, can be, e.g., a sensory experience from which that belief sprang. In that case, the belief would be justified by the experience (or by its origin from the experience) only if S justifiably believes that the belief sprang from that experience and that this origin is sufficient for justification. On this version the perspective determines justification by determining what can justify what; but

it allows items outside the perspective (items other than justified beliefs of that subject) to function as justifiers.⁷

Here and elsewhere in the paper the following distinction will be useful. A belief is *mediately (indirectly)* justified provided it is justified by virtue of its relations to other justified beliefs of the subject that provide adequate support for it. In such cases the belief is justified by the *mediation* of those other beliefs. If it is justified in any other way it will be said to be *immediately (directly)* justified. In terms of this distinction, the view embodied in the first quotation from Bonjour rules out *purely* immediate justification, justification by something other than other justified beliefs of the subject *alone*, since it holds that an experience can justify a belief only if the subject has certain justified beliefs about the experience and its relation to the belief; but it is hospitable to mixed justification, in which both other justified beliefs *and* something else are required for justification.⁸ However there is or can be a version of PI that is more radically opposed to immediate justification, one that would “perspectivize” justifiers more thoroughly, by holding that only justified beliefs can *be* justifiers. On this version what justifies a perceptual belief is not the experience itself, or actual origin from the experience, but the justified belief that the experience has occurred or that the belief originated from it.

We have made the distinction between these versions hang on what is allowed to count as “a justifier”. In the perceptual case both versions require a justified belief that the relevant experience occurs, but they differ as to whether the experience itself can function in a justifying role. But this might be thought a trivial verbal difference, having to do only with where we draw the line between what is doing the justifying, and the conditions under which it is enabled to do so. What difference does it make where that line is drawn? On both views both “the justifier” and “the conditions that must obtain if it is to be a justifier” figure essentially in the conditions that are necessary for the belief in question to be justified. Why does it matter how we divide that set of conditions into what *does* the justifying and what *enables* it to do that justifying?

I agree that the division is not of any great importance. Nevertheless there is an important difference between the versions. For Bonjour’s version, in allowing the experience itself into the necessary conditions for justification, under whatever rubric, is imposing a condition for the justification of the perceptual belief over and above those imposed by the more radical alternative. Put it this way. Both versions alike hold that S is justified in believing that p (that there is a tree in front of one) only if S is justified in believing that S has experience E. But Bonjour imposes the additional requirement that S *have* the experience; i.e., he requires that the supporting belief be *true*. And this can be seen to mark a decisive superiority of the more radical alternative. We are

dealing with a case in which S’s belief that s/he has experience E provides him/her with an adequate reason for the perceptual belief. (If more justified beliefs on S’s part, about normality or other background conditions, are required for this, let them be included also.) Otherwise the case would fall short of justification by reason of the insufficiency of the alleged ground and we would never get to the problems raised by the internalism-externalism distinction. But if I do justifiably believe that I am having E, and if that constitutes a sufficient reason for my supposing that p, that is surely enough for my being justified in believing that p. To require that my supporting beliefs be *true* might be appropriate if we were laying down requirements for knowledge, but it is clearly too strong a requirement for justification. If, e.g., I am justified, to as high degree as you like, in supposing that my car is in my garage, then I am surely *justified* in denying that it is parked in front of the bank, even if, unbeknownst to me, someone had removed it from my garage and parked it in front of the bank. Thus Bonjour’s version represents something of an overkill.⁹ Let’s codify the preferred version.

- (3) The only thing that can justify S’s belief that p is some other justified beliefs of S.

Next let’s note a respect in which (3) needs broadening. Recall the important notion of *prima facie* justification. One is *prima facie* justified in believing that p provided that one is so situated that one will be (unqualifiedly, all things considered, *ultima facie*) justified in believing that p, provided there are no sufficient “overriding”¹⁰ considerations. Thus in a normal perceptual situation in which I take myself to see a tree in front of me I am thereby *prima facie* justified in believing that there is a tree in front of me; but this justification can be overridden by abnormalities in the situation, e.g., sensory malfunctioning of various sorts. Now consider what a PI internalist should say about the conditions under which a *prima facie* justification is overthrown. Does the mere existence of a sufficiently serious malfunctioning suffice? Or would the subject have to know or be justified in believing that this was the case? Clearly it is the second alternative that is in the spirit of PI. Just as the mere fact that a belief was produced in a highly reliable manner cannot justify it, so the mere fact that a belief was generated in an unreliable fashion cannot serve to discredit the belief. In both cases justification, or the lack thereof, depends on how the situation appears within my perspective, i.e., on what I know or justifiably believe about it. If and only if I have sufficient reason to think there to be something fishy about this case of perception, will *prima facie* justification be overthrown. And, indeed, most epistemologists have taken this line about what overrides *prima facie* justification, even where they haven’t also

accepted (3) as a constraint on justification.¹¹ Thus we should add overrides to the scope of (3). In the interest of concise formulation let us introduce the term "epistemizer" to range over anything that affects the justification of a belief, positively or negatively. We can then reformulate (3) as:

- (4) The only thing that can epistemize S's belief that p is some other justified belief(s) of S.

Now we are in a position to return to the choice between knowledge, belief, and justified belief in the specification of the subject's perspective. To deal with this properly we must note that (4) places severe restrictions on a theory of justification by implying that only *mediate* justification is available. Let's call any theory of justification that recognizes only mediate justification a discursive theory. Whether there are other varieties depends on how narrowly the boundaries of coherentism are drawn, and there is wide variation on this. For the present let's think of coherentism widely, as ranging over any discursive theory.

Next let's distinguish between "positive" and "negative" coherence theories. John Pollock introduced the distinction as follows.

There are two kinds of coherence theories. On the one hand, there are coherence theories which take all propositions to be *prima facie* justified. According to those theories, if one believes a proposition, P, one is automatically justified in doing so unless one has a reason for rejecting the belief. According to theories of this sort, reasons function primarily in a negative way, leading us to reject beliefs but not being required for the justified acquisition of belief. Let us call these negative coherence theories. The other kind of coherence theory (a positive coherence theory) demands positive support for all beliefs.¹²

In other words, on a positive coherence theory a belief is justified only if it stands in the right relation to justifiers. On a negative coherence theory a belief is justified unless it stands in the wrong relation with overrides. What makes them both coherence theories is that in both cases the epistemizers must be drawn from the subject's propositional attitudes.

Now let's go back to the various sorts of propositional attitudes that might be supposed by PI to make up the subject's perspective: beliefs, justified beliefs, knowledge. In formulations (2), (3), and (4) we chose *justified belief* without explaining or justifying that choice. I now turn to that task.

First what about the decision between knowledge and justified belief? Here the point is that the more modest constraint is called for. Suppose that I am

justified in believing that my car is in the garage, since I left it there this morning and have been away from the house since, no one else has a key to the house or garage, and the neighborhood is remarkably free of crime. In the afternoon I see a car that looks like mine in the parking lot of a bank, but believe that it isn't mine, on the grounds of my car's being in my garage. Suppose further that my car has been stolen and this is my car, so I didn't *know* that my car was in the garage even though I was justified in believing this. I am surely justified in believing that the car in the parking lot is not mine, even though the basis for this belief is something I am justified in believing but do not know. Cases like this indicate that it is sufficient for a belief to be a justifier that it be justified; it is not also required that it count as knowledge.

But what about the alternative between any beliefs on the one hand, and only justified beliefs on the other? It may seem that we can settle this issue in the same way. Suppose I merely believe that my car is in the garage, just because that is where I normally expect it to be when I don't have it with me; but I am not justified in believing this. On the contrary, I took it to a repair shop to be worked on yesterday; when I believe that it is in my garage I have temporarily forgotten about this incident, even though I am quite capable of remembering it and have been remembering it except for this temporary lapse. Again I take my car's being in my garage as a reason for supposing that the car I see in the bank parking lot is not mine. Here it is quite clear that I am not justified in this latter belief by virtue of basing it on an unjustified belief. On the contrary, the fact that I am quite unjustified in supposing my car to be in my garage shows that I don't become justified in some further belief by virtue of basing it on that belief. More generally, it seems that beliefs cannot acquire justification by being brought into relation with unjustified beliefs. One belief cannot "transfer" to another belief a justification it does not possess.

This last argument is, I believe, conclusive for what we might call "local" mediate justification, justification of a particular belief by the evidential or other logical relations in which it stands to one, or a few, other beliefs. Justification can be transferred "locally" only by beliefs that already have it. But the more common sort of discursive theory is a "holistic" coherence theory, one which takes a given belief to be justified, at least in the last analysis, not by its relations to a very few other beliefs "in the vicinity", but by the way in which it fits into some very large system of beliefs. Since the term "coherence theory" derives from the idea that a belief is justified if and only if it "coheres" with such a total system, it will be most natural to restrict the term "coherence theory" to holistic theories. The obvious choice for the system with which a belief must cohere in order to be justified is the totality of the subject's

current beliefs. Thus on the most usual sort of coherence theory the subject's "perspective" by reference to which the justification of any particular belief is to be assessed consists of the subject's beliefs, without any further restriction to justified beliefs. Indeed, there could not be such a restriction. For on the kind of (pure) coherence theory we are now considering a belief is or is not justified just by its relations to the whole of the subject's beliefs. Apart from that coherence with *all* the subject's beliefs there are no justified beliefs to serve as a reference class. Hence by the time the totality of beliefs has been segregated into justified and unjustified it is too late to use the former class as a touchstone to determine whether a given belief is justified. That determination has already been made. Of course, if at a future time the subject has some new beliefs, we can at that time assess their justificatory status, and this determination will be made after the earlier demarcation of the justified from the unjustified beliefs. But that doesn't change the verdict. At that future time, by the terms of the theory, a given belief (new or old) is justified solely on the basis of its coherence with the total set of beliefs the subject has at that time. And so for a pure coherence theory PI should be formulated as follows.

- (5) Only the total set of S's beliefs at t can function as an epistemizer at t.¹³

Since I find pure coherence theories quite unsatisfactory for a variety of reasons, I might seek to rule out (5) on those grounds. But in this paper I did not want to get into *substantive* epistemological issues like those concerning the opposition between foundationalism and coherentism. This paper is designed to be restricted to meta-epistemological issues concerning basic epistemological concepts, their explication, interrelations, and suitability for one or another purpose. Thus I shall just point out that the internalism-externalism dispute is mostly carried on by thinkers who believe in local mediate justification. Hence we will ensure maximum contact with that debate if we focus on (4) rather than (5) in the ensuing discussion.

One more point must be laid on the table before we turn to the consideration of what can be said in support of PI. Go back to the initial quotation from Bonjour; we have not yet squeezed it dry. There Bonjour requires for the justification of S's putatively basic belief that S justifiably believe not only that the belief have the "feature" in question but also *that beliefs having that feature are likely to be true*. When we come to the main argument for PI we will see the rationale for this additional higher level requirement. For the moment we need only note its general character. It is clear that Bonjour imposes *this* requirement just because he takes truth conducivity to be required for, as we might say, *justificatory efficacy*. Earlier in the essay he had written that "the dis-

tinguishing characteristic" of epistemic justification is "its internal relation to the cognitive goal of truth" (p. 54). Elsewhere Bonjour has laid it down that it is essential to a justifier to be "truth conducive".¹⁴ Thus this additional requirement is really a requirement to the effect that the subject be justified in supposing not only that the putative justifier obtains but also that it be efficacious, that it have what it takes to justify the belief. But he can't come right out and say that. Consider his situation if he were to try. Formulate the additional requirement as: *S justifiably believes that the possession of that feature suffices to justify the belief*. But Bonjour is committed to deny this; his specific contention is that no feature of a belief can be sufficient to justify the belief; the subject must also have certain justified beliefs about that feature.¹⁵ Then how about requiring that the possession of that feature is part of what confers justification on this belief? But we want the requirement to be more specific than that. The two justified beliefs are also part of what confers justification on the belief in this situation, but a different part. We want to specify what part the feature is contributing to the justification. That is what Bonjour is attempting to do with his requirement that S be justified in believing that the feature is probabilifying, that by virtue of having this feature the belief is likely to be true. That will do the job, on the assumption that probabilification is just what it takes for justificatory efficacy. But this is controversial. In fact, other internalists have been in the forefront of denying just this claim.¹⁶ Thus it appears that if we are to give an adequate formulation of this higher level requirement we must commit ourselves to some highly controversial assumption as to what is required for justification, some highly controversial assumption in substantive epistemology.

Fortunately there is a coward's way out, since we are working with (4), which restricts us to purely mediate justification, rather than with Bonjour's versions. On (4) the only justifiers are other justified beliefs of the same subject. Hence the way in which any justifier has to be related to a belief in order to do its job is to provide "adequate support" or "adequate evidence", to be an "adequate reason". No doubt, it is both obscure and controversial what is required for one belief (or the propositional content thereof) to constitute an adequate reason or to provide adequate support for another. But leaving all this aside, and taking cover behind the criterion-neutral term "adequate", we can put the additional, higher-level requirement just by saying that S must justifiably believe that the justifying belief(s) provide adequate support for the justified belief.¹⁷ Tacking this on to our canonical formulation we get:

- (6) Only S's justified beliefs can epistemize S's belief that p, and then only if S justifiably believes that the other justified beliefs in question

provide adequate support for p (or for something else, in the case of overriders).

II

We have now explicated PI sufficiently to consider what can be said in its favor. That consideration will lead to further refinements. First let's consider what defense is offered by BonJour in "Externalist Theories of Empirical Knowledge." The main effort there is devoted to an attack on reliability theories, utilizing an example of alleged clairvoyance. It is stipulated that the subject has a reliable capacity for determining the disposition of distant objects on no apparent basis. BonJour first argues that if the person has adequate reason for supposing that a belief thus formed is false, or that her clairvoyance is not reliable, then she is not justified in the clairvoyant beliefs, even though they are formed reliably. But, as BonJour acknowledges, this shows only that the subject's justified beliefs do have a bearing on what other beliefs are justified, not that they are the only thing that have this bearing. Next, he more boldly argues that in the case in which the subject has no reasons for or against the reliability of her powers or the truth of the belief (whether or not she believes that the powers are reliable), she is not justified in holding the beliefs, however reliable her clairvoyant powers are in fact. However these "arguments" simply consist in BonJour's displaying his intuitions in opposition to those of his opponent. A couple of quotations will give the flavor.

We are now face-to-face with the fundamental – and seemingly obvious – intuitive problem with externalism: *why* should the mere fact that such an external relation (the reliability of the faculty) obtains mean that Norman's belief is epistemically justified, when the relation in question is entirely outside his ken?

One reason why externalism may seem initially plausible is that if the external relation in question genuinely obtains, then Norman will in fact not go wrong in accepting the belief, and it is, *in a sense*, not an accident that this is so. But how is this supposed to justify Norman's belief? From his objective perspective, it *is* an accident that the belief is true.¹⁸

This is more like an appeal to PI than a support for that restriction. There are, as we shall see, some germs of a more substantial argument in BonJour, but they will need developing.

Nor are we helped by a rather common argument for PI that stems from a confusion between the *activity* of justifying a belief and the *state* of a belief's being justified. Here is a good sample.

In whatever way a man might attempt to justify his beliefs, whether to himself or to another, he must always appeal to some belief. There is nothing other than one's belief to which one can appeal in the justification of belief. There is no exit from the circle of one's beliefs.¹⁹

Of course, if I am to carry out the *activity* of justifying a belief I must provide an argument for it; I must say something as to why one should suppose it to be true. And to do this I must employ other beliefs of mine. In saying what reasons there are for supposing that p I am expressing other beliefs of mine, and contextually implying that I am justified in accepting them. But this all has to do with the activity of *justifying* a belief, *showing* it to be justified. From the fact that I can *justify* a belief only by relating it to other beliefs that constitute a support, it does not follow that a belief can *be justified* only by its relations to other beliefs. Analogously, from the fact that I cannot justify my expenses without saying something in support of my having made them, it does not follow that my expenses cannot *be justified* unless I say something in support of my having made them. Indeed, we all have innumerable beliefs that are commonly taken to be justified but for which we never so much as attempt to produce reasons. It might be argued with some show of plausibility that one can be justified in believing that p only if it is *possible* for one to justify that belief; but I cannot imagine any remotely plausible argument for the thesis that I can be justified in believing that p only if I *have justified* that belief. Hence the point made by Lehrer about justifying leaves completely intact the possibility that one might *be justified* in a belief by something other than one's other beliefs.

We will have to make the same judgment on an analogous argument from what is involved in *deciding* what to believe. Here is a version by Pollock.

In deciding what to believe, we have only our own beliefs to which we can appeal. If our beliefs mutually support our believing P, then it would be irrational for us not to believe P and hence belief in P is justified. There is no way that one can break out of the circle of his own beliefs.²⁰

Again, even if this shows that I can have no basis other than my own beliefs for a *decision* as to what to believe, it falls far short of showing that nothing

can *justify* a belief except other beliefs. For there is no reason to suppose that the only justified beliefs are those the subject *decided* to adopt.

Even though, as will appear in the fullness of time, I am no advocate of PI, I feel that I can improve on the recommendations for that view that can be found in the writings of its supporters. Here is what I take to be the strongest argument for it. I have gleaned the basic idea for this line of argument from various sources, but the development of it is my own.²¹

"First let's note that the fact that q can enter into the justification for S's believing that p only in the guise of S's being justified in believing that q. Consider the popular idea that what justifies me in beliefs about my own current conscious states is that such beliefs are infallible, i.e., are such that I couldn't mistakenly form such a belief. But how could that fact justify those beliefs unless I were cognizant of the infallibility? If I am unaware of their infallibility, and they have no other justification, am I not proceeding *irresponsibly* in forming such beliefs? Just as the mere fact that X is about to attack me will not justify my striking X unless I have good reason to suppose that he is about to attack me, so the mere fact that current feeling beliefs are infallible can't justify me in accepting them unless I at least have good reason to regard them as infallible. *Pari passu*, the mere fact that I am being appeared to *treely* cannot render me justified in believing that there is a tree in front of me, unless I am justified in believing that I am being appeared to *treely*. If I am unaware of the existence of the warrant conferring fact then, for me, it is just as if it did not exist. How can a fact of which I take no account whatever have any bearing on what it is *permissible* for me to do, in the way of action or of belief? Thus it would seem that my being justified in believing that q is at least a *necessary* condition of q's playing a role in justifying my belief that p.

But it is also a sufficient condition. Provided I am justified in believing that beliefs about current feelings of the subject are infallible, what more could be required to legitimate those beliefs? Even if they are not in fact infallible, how can that prevent its being *permissible* for me to accept them? If, so far as I can tell, there are facts that strongly support the supposition that p, then surely it is *all right* for me to give my assent to p. What more could be *demanded* of me? I have done all I can. What the actual facts are over and above what I am most justified in believing is something I cannot be held *responsible* for. Once I have marshalled all the cognitive resources available to me to determine the matter I have, in my body of justified beliefs, the closest approximation I can make to the actual facts. That is the best I have to go on, and it would be quite unreasonable to suggest that I *ought* to be going on something else instead. What I am justified in believing provides sufficient as well as necessary conditions for the justification of further beliefs."²²

How does this line of argument go beyond simply displaying internalist (PI) intuitions? It does so by grounding those intuitions in a particular conception of justification, one that makes epistemic justification a matter of the subject's normative situation, a matter of how the subject's believing that p stands *vis-à-vis* relevant intellectual norms, standards, obligations, duties, and the like. If S's believing that p is *not* in contravention of relevant intellectual obligations, then it is *permissible* for him to believe that p, he cannot be rightly *blamed* for doing so, it is *all right* for him to hold that belief, he is *in the clear* in so believing. Let's call this a "deontological" conception of epistemic justification. The argument just presented exhibits the PI constraint as flowing from *what justification is*, as thus conceived. Since whether I am justified in believing that p depends on whether I could rightfully be blamed or held to account for so believing, then what is crucial for whether I am justified is the way the relevant facts appear from my perspective; justification depends on what the relevant facts are like, *so far as I can tell*. For that is what is crucial for whether I can be blamed for my belief. If and only if my belief is adequately supported *so far as I can tell*, I cannot be blamed for the belief.²³

Elsewhere I have explored the deontological conception and contrasted it with the very different "strong position" (SP) conception, as well as distinguishing various versions of each.²⁴ Roughly speaking, to be SP justified in believing that p is to believe that p in such a way as to be in a strong position thereby to attain the truth and avoid error. It is to believe that p in a "truth conducive" way. It is for one's belief to have been formed in such a way or on such a basis that one is thereby likely to be believing correctly. Note that each conception omits the crucial emphasis of the other, thereby implicitly denying it to be necessary for justification. Freedom from blameworthiness, being in the clear as far as one's intellectual duties are concerned, is totally ignored by the "strong position" theorist. So long as one forms one's belief in a way that is well calculated to get the truth, it is of no concern how well one is carrying out intellectual duties. Conversely, the deontologist has nothing to say about truth conductivity.²⁵ So long as I am not violating any intellectual duties I am "in the clear" in believing that p, whatever my chances for truth. This is not to say that each side denies the importance of what is crucial for the other. The deontologist need not be indifferent to the truth, nor need the "strong position" theorist be uninterested in intellectual duties. But they differ on how these admittedly important matters relate to epistemic justification.

To get a properly rounded picture we should also note a way in which truth conductivity does typically enter into deontological theories of justification. Even though truth-conductivity does not enter into the meaning of "justified" for the deontologist he is likely to give it a prominent place when

he comes to spell out the content of our most important intellectual obligations. Such theorists typically hold that our basic intellectual obligation is to so conduct our cognitive activities as to maximize the chances of believing the true and avoiding believing the false.²⁶ Thus even though one may be deontologically justified without thereby being in a favorable position to get the truth, if our basic intellectual obligation is to maximize truth and minimize falsity one cannot be deontologically justified in a belief unless one is believing in such a way that, so far as one can tell, is well calculated to reach the truth.

Now we can see that just as the deontological conception supports a PI restriction, so an SP conception supports its denial. It is obviously not conceptually necessary that one comes to believe that *p* in a truth conducive way only if that belief is well supported by other justified beliefs of the subject. It is clearly possible that there are ways of being in a strong position in one's beliefs other than by basing those beliefs on other justified beliefs. Plausible examples of such other ways are not far to seek. Perceptual beliefs about the physical environment, e.g., that the lilies are blooming in the garden, are based on the subject's sensory experience, on the way in which things sensorily appear to one. Furthermore let's make the plausible supposition that one does not typically form beliefs about how one is being sensorily appeared to; the sensory appearance directly gives rise to the belief about the environment. It is not that one says to oneself, even rapidly, implicitly, or below the level of consciousness, "I am having a visual experience of such and such a sort; therefore the lilies are blooming in the garden". No such inference typically takes place, for the premises for such inferences are rarely made objects of belief. Finally, let's make the plausible assumption that our perceptual belief forming mechanisms are generally reliable, at least for the sorts of perceptual beliefs we typically form, in the sorts of situations we typically encounter. Granting all this, perceptual belief formation constitutes massive support for the thesis that one can form beliefs in a reliable, truth conducive, manner, without basing them on other justified beliefs.

Beliefs about ones current conscious states provide even stronger support. It is very plausible to suppose that we have a highly reliable (some would even say infallible) mechanism for the formation of such beliefs. And yet it would be extremely implausible to suppose that these beliefs are formed or held on the basis of reasons. What would such reasons be? It may be suggested that my reason for supposing that I feel sleepy at the moment is that I do believe this and that such beliefs are infallible. But many persons who form such beliefs do not even have the relevant concept of infallibility, much less typically believe that such beliefs are infallible whenever they come to believe

such things. Once again we have reason to suppose that beliefs can satisfy the SP conception of justification without satisfying the PI constraint on justification.

Next let's note that the argument we have given for PI supports both the lower level and the higher level requirement laid down in (6). The "lower level requirement" is that the belief that *p* be adequately supported by other justified beliefs of the subject, and the "higher level requirement" is that the subject justifiably believe that these other justified beliefs provide adequate support for the belief that *p*. We have been emphasizing the way in which the argument establishes the lower level requirement, but it also lends powerful support to the higher level requirement. For suppose that my belief that *p* is based on other justified beliefs of mine and, let's suppose, these other justified beliefs provide adequate support for the belief that *p*. But suppose further that I do not justifiably believe that these other beliefs do provide adequate support. In that case, so far as I can tell, I do not have within my perspective adequate support for *p*. Would I not be proceeding irresponsibly in adopting the belief that *p*? Couldn't I properly be held accountable for a violation of intellectual obligations in giving my assent to *p* under those conditions? Therefore if I am to be in the clear in believing that *p*, the belief must not only be based on other justified beliefs of mine; I must also be justified in supposing those beliefs to provide sufficient support for the belief that *p*.

That shows that the higher level belief is necessary for justification. We can now proceed to argue that it, together with the lower level requirement, is sufficient. The crucial question here is whether it is also necessary for justification that the other justified beliefs do in fact provide adequate support, that their propositional contents are indeed so related as to make the one an adequate reason for the other. A consideration of conditions of blame, being in the clear, etc., will support a negative answer. For if, going on what I know or justifiably believe about the world, it is clear to me that other justified beliefs of mine adequately support the belief that *p*, what more could be required of me? Even if I am mistaken in that judgment I made it in the light of the best considerations available to me. I can't be held to blame if I proceed in the light of the best reading of the facts of which I am capable. Hence a *justified belief* that I have adequate support is all that can rightfully be imposed in the way of a higher level requirement.

Now that we have a two level PI internalism-externalism contrast, there is the possibility of being an internalist on one level and an externalist on another. The two parties disagree both over what can be a justifier and over that by virtue of which a particular item justifies a particular belief. A

particularly live possibility of a compromise is an internalism as to what can justify and an externalism as to what enables it to justify. One could be a PI internalist about justifiers by virtue of recognizing only mediate justification, but insist that my belief that *p* is justified by its relations to my belief that *q* if and only if *q* does in fact provide adequate support for *p*. At the end of the paper we shall advocate a similar mediating position, though the internalist component will not be the PI brand.

Now let's consider a way in which what is supported by our argument for PI differs from the formulation of PI with which we have been working. We have represented the deontologist as maintaining that whether *S* is justified in believing that *p* is solely a function of what other justified beliefs *S* has. But that cannot be the whole story. Consider a case in which, although the sum total of the justified beliefs I actually possess provide an adequate basis for the belief that *p*, that would not have been the case had I been conducting myself properly. If I had looked into the matter as thoroughly as I should have, I would be in possession of effective *overrides* for my evidence for *p*, and my total body of evidence would not have given sufficient support for the belief that *p*. Here the belief that *p* is adequately supported by the perspective on the world that I actually have, and I justifiably believe that it is; but nevertheless I am not in the clear in believing that *p*, not justified in the deontological sense.

These considerations show that PI must be modified if it is to be supported by a deontological conception of justification. It must include a codicil to the effect that overrides that the subject does not possess, but would have possessed had s/he been conducting him/herself as s/he should have been, also can serve to epistemize beliefs.²⁷ PI now becomes:

- (7) Only *S*'s justified beliefs can epistemize *S*'s belief that *p*, and then only if *S* justifiably believes that those other justified beliefs provide adequate support for *S*'s belief that *p*; but overrides that *S* should have had but didn't can cancel out justification provided by the preceding.²⁸

Going back once more to our argument for PI, I now wish to point out that it utilizes a special form of a deontological conception of justification that is limited in ways that render it either totally inapplicable, or at least severely limited in application.

First it utilizes a concept of justification that assumes beliefs to be under direct voluntary control. The argument takes it that one is justified in believing that *p* if and only if one is not to blame for believing that *p*, if and only if in that situation this was a belief that one was permitted to choose. All this talk has

application only if one has direct voluntary control over whether one believes that *p* at a given moment. If I lack such control, if I cannot believe or refrain from believing that *p* at will, then it is futile to discuss whether I am permitted to believe that *p* at *t* or whether I would be irresponsible in choosing to believe that *p* at *t*. And it seems that we just don't have any such control, at least not in general. For the most part my beliefs are formed willy-nilly. When I see a truck coming down the street I am hardly at liberty either to believe that a truck is coming down the street or refrain from that belief. Even if there are special cases, such as moral or religious beliefs, where we do have pin point voluntary control (and even this may be doubted), it is clear that for the most part we lack such powers.²⁹

Not only does the argument in question presuppose direct voluntary control of belief. It considers the requirements for justification only for those beliefs that are acquired by an explicit, deliberate choice. For it arrives at the PI constraint by pointing out that only what I am cognizant of can be taken account of in my *decision* as to whether to believe that *p*. "If I am unaware of their infallibility, . . . am I not proceeding irresponsibly in forming such beliefs?" "If, so far as I can tell, there are facts that strongly support the supposition that *p*, then surely it is all right for me to give my assent to *p*." But this fact, that only what I am cognizant of can affect the permissibility of my choice, will imply a *general* constraint on justification only if justification is confined to beliefs that are *chosen* by a deliberate voluntary act. But even if beliefs are *subject to* direct voluntary control that control need not always be exercised. One can hold that it is always in principle possible to choose whether to believe a given proposition without thereby being committed to the grossly implausible supposition that all our beliefs are in fact acquired by an explicit choice. Even overt actions that are uncontroversially under voluntary control, such as tying one's shoelaces, can be, and often are, performed habitually. Likewise, even if beliefs are as subject to direct voluntary control as tying one's shoelaces, beliefs are often acquired willy-nilly. Even if we are always capable of voluntarily controlling our perceptual belief formation, e.g., we rarely do so. Hence a concept of epistemic justification that is confined to beliefs acquired by deliberate choice covers only a small part of the territory.

Third, it follows from the point just made that the argument utilizes a concept of justification that evaluates a belief solely in terms of its original acquisition, for the argument has to do with what can determine the permissibility of the *choice* of a belief. But it is often noted by epistemologists that the epistemic status of a belief may change after its acquisition, as the subject comes to acquire or lose support for it. Suppose that after coming to believe that Susie is quitting her job, on the basis of no evidence worthy of the name

and hence unjustifiably, I come into possession of adequate evidence for this supposition; let us further suppose that this new evidence now functions as the basis for my belief. In this case my belief comes to be justified *after* its acquisition. Thus a concept of justifiably *acquired* belief is at best only a part of an adequate concept of justified belief.

To be sure, it is not difficult to modify this very restrictive concept, so as to make it more generally applicable. Let's begin by showing how the direct voluntary control assumption can be dropped. It is uncontroversial that our beliefs are under *indirect* voluntary control, or at least subject to influence from our voluntary actions. Even if I can't effectively decide at this moment to stop believing that Reagan is inept, I could embark on a regimen that is designed to improve my assessment of Reagan, and it might even succeed in time. With this possibility of indirect influence in mind, we can reconstrue "intellectual obligations" so that they no longer attach to believings and abstentions therefrom, but to actions that are designed to influence our believings and abstentions. Reinterpreted in this way the argument would be that whether we are justified in believing that *p* at *t* would depend on whether prior to *t* we had done what could reasonably be expected of us to influence that belief. The difference between these two understandings may be illustrated as follows. Suppose that my belief that there is life outside our solar system is inadequately supported by the totality of my justified beliefs. On the direct voluntary control interpretation I have an effective choice, whenever I consider the matter, as to whether to keep believing that or not. It is my duty to refrain from believing it since it is not adequately supported by my "perspective"; since I continue to believe it in defiance of my duty I am doing something that is not permitted; my belief is not justified. But the matter sorts out differently on the "indirect voluntary control" construal. It is recognized that I lack the capacity to discard that belief at will; at most I have the ability to make various moves that increase the chances of the belief's being abandoned. Hence so long as I am doing as much along that line as could reasonably be expected of me, I can't be faulted for continuing to have the belief; and so it is justified. On either of these interpretations, whether my belief is justified is a function of how things appear in my perspective rather than of how they are in actual fact. So long as life outside the solar system is improbably relative to what I am justified in believing, then my belief is unjustified unless (on the indirect control version) my best efforts have failed to dislodge it.

Next consider how we can lift the other restrictions. We can confine this discussion to the direct control version, since on the indirect control version there was no reason to impose them in the first place. Let's first take

the restriction to explicitly chosen beliefs. On the direct control version we can say that the belief is justified provided that it was acquired on such a basis that if the agent had chosen to adopt the belief on that basis he could not have been blamed for doing so. In other words, where the belief, or its furtherance, was not explicitly chosen we can evaluate it, on the deontological conception, by considering whether its basis is such that if it or its furtherance was chosen on that basis the agent would have been in the clear in so choosing.

Now let's see how to lift the restriction to the original acquisition of the belief, and extend the concept to the evaluation of one's continuing to believe that *p* at times after its original acquisition. Once again the crucial move is to consider what would be the case if we were to make a choice that we did not in fact make. For one thing, we can consider what the judgment would be on my coming to believe this if the belief were voluntarily adopted on the basis of this evidence I possess at present (and analogously for the indirect control version). Or, closer to home, we could consider the possibility that I should now explicitly raise the question of whether to retain the belief, in the light of the evidence I now possess, and should come to a decision to retain it. In that case would I be in the clear in making that decision? If so I am now justified in retaining the belief.

It is time to take our bearings with respect to these increasingly proliferating variations in a deontological concept of justification. To keep complexity within manageable bounds, I shall formulate a version that is designed to take care both of habitually formed beliefs and post-acquisition influences on justifications. I shall formulate this both in a direct control and an indirect control version.

- (8) Direct control version. One is justified in believing that *p* at *t* if and only if either (a) in choosing at *t* to adopt or retain the belief that *p* one was not violating any intellectual obligations, or (b) one's belief that *p* at *t* has such a basis that if one were to decide, in the light of that basis, to retain one's belief that *p*, one would not be violating any intellectual obligations in so doing.
- (9) Indirect control version. One is justified in believing that *p* at *t* if and only if one's believing that *p* at *t* does not stem from any violations of intellectual obligations.

Thus it is not difficult to concoct distinctively deontological conceptions of justification that avoid the severe limitations of the concept employed by the

argument for PI. But what sort of argument for PI can be constructed on the basis of these alternative conceptions?

The first point is that no case at all can be made for PI on the basis of the indirect control version. According to (9), justification is a function of certain features of the causal history of the belief. Was that history such that if the subject had lived up to his/her intellectual obligations in the past then s/he would not have believed that *p*? This is not a "perspectival" matter. The justified beliefs of the subject do not play any crucial role in determining whether or not that condition was satisfied. It is matter of what actually went on, rather than a matter of how what went on is represented in the subject's viewpoint. Thus (9) supports an externalist position on justification; at least it supports the externalist contrast to PI. Of course we could try to "perspectivize" (9). Any condition for anything that is in terms of what the facts actually are can receive a "perspectival" modification, transforming it into a condition that the facts be represented in a certain way in the subject's perspective. So modified (9) would become:

- (10) S is justified in believing that *p* if and only if S's belief that *p* did not, so far as S can tell, stem from S's violations of intellectual obligations.

But (10) is wildly permissive. We rarely have reason to think that one of our beliefs stems from intellectual transgressions. To know about the causal history of beliefs takes research, and we rarely engage in such research. Hence we have very few beliefs about the causal history of our beliefs. And so practically all beliefs, no matter how shoddy or disreputable, will be justified on this criterion. The prospects for support for PI from an indirect control version of a deontological conception are vanishingly small.

Things do not look much rosier from the perspective of (8). According to (8) a belief can be justified on the basis of anything whatsoever, not just other justified beliefs of the subject, provided that one would be in the clear, vis-à-vis one's intellectual obligations if one were to consider whether to retain the belief in the light of that basis. If one were to engage in such a consideration one would, of course, be choosing to retain the belief on the basis of other justified beliefs, in particular the belief that that basis obtains. That is the situation envisaged by the restrictive concept employed in the original argument for PI. But the extended concept differs from that precisely by not making the actual obtaining of such a situation necessary for justification. It recognizes that a belief can be justified even if one never does make any decision with respect to it on the basis of what one justifiably believes about its basis. Hence on this modified deontological concept a belief could be justified by being

based on some experience, even if the subject in fact has no beliefs about that experience. What is supported by (8) is a denial rather than an affirmation of the PI constraint.

Thus it appears that we have a significant argument for PI only if we utilise a concept of justification that cannot be seriously defended as generally applicable, a concept according to which the justification of beliefs is solely a matter of whether a belief is *chosen* in such a way that this choice does not involve any dereliction of intellectual duty. But we cannot seriously suppose that justified beliefs are restricted to those that are *chosen* in that way, even if some are. Insofar as we are working with an even minimally defensible concept of justification, the argument for PI dissipates.

When we consider the higher level requirement embodied in (7), things look even worse. (7) implies that I will be justified in believing that *p* on the basis of my justified belief that *q* only if I am justified in supposing that the latter belief provides adequate support for the former. One reason this darkens the prospects for PI is that it is doubtful that we satisfy that condition very often. Just how often it is satisfied depends on what it takes to be justified in beliefs like that, and that is not at all clear. One thing that is clear for the PI advocate, however, is that to be justified such a belief will have to be mediately justified, since that is the only kind of justification PI recognizes. We will have to have sufficient reasons for supposing that *the belief that q adequately supports the belief that p* if we are to be justified in that higher level belief. How often do we have such reasons? Not very often, I would suggest. Perhaps the following will suffice to indicate the difficulties. Consider perceptual beliefs. If my perceptual belief that it is raining outside is to be mediately justified, this will presumably be on the basis of a justified belief that I am having certain visual experiences, plus perhaps (depending on the requirements we adopt) justified beliefs about the normality of the situation.³⁰ Now to have adequate reasons for supposing that reasons like that are sufficient support for a perceptual belief about one's environment is to be in the position that many great philosophers have labored to get themselves into when they have wrestled with the problem of how to infer facts about the external world from facts about the sensory experiences of the individual percipient. And even if some philosophers have solved that problem, which I am strongly inclined to deny, it is quite clear that the overwhelming majority of the population is not in possession of any such solution. For a second illustration, consider the point that in order for some non-deductive evidence to be adequate support for a given belief (so that this latter belief is justifiably held) there must be no other justified beliefs of mine that serve to defeat the *prima facie* support provided by the first-mentioned evidence. Suppose that my reason for supposing that Ray will be in his office today is that today is Wednesday and Ray has a fixed

disposition to work in his office on Wednesdays. However I have temporarily forgotten that Ray told me last week that he will be out of town on Wednesday of this week. When that justified belief of mine is added to the picture the total evidence no longer adequately supports the supposition that Ray will be in his office today. This means that I can be justified in supposing that my belief that *q* renders my belief that *p* justified only if I am justified in supposing that there is nothing else I am justified in believing such that when that is added to *q* the conjunction does not adequately support *p*. And it is difficult to be justified in any claim concerning what is or is not present in the totality of ones justified beliefs.

Thus it is dubious that the higher level requirement of PI is very widely satisfied. If that is required for justification not many people are justified in many beliefs. But there is an even more serious difficulty with the requirement. It engenders an infinite regress. If in order to be justified in believing that *p*, I must be justified in believing that my reason, *q*, adequately supports *p*, the justification of this later belief requires the justification of a still higher level belief. That is, if *r* is my reason for supposing that *q* adequately supports *p*, I can be justified in supposing that *q* adequately supports *p* only if I am justified in supposing that *r* adequately supports **q adequately supports p**. And my justification for this last belief includes my being justified in a still higher level belief about adequate support. Given PI, I cannot be justified in any belief without simultaneously being justified in all the members of an infinite hierarchy of beliefs of ever ascending level.

Let's make sure we fully appreciate the character of this difficulty. The view that all justification is mediate itself gives rise to a much more widely advertised regress, this one stemming from the lower level requirement that a given belief can be justified only by its relation to another justified belief. The same is true of the justification of this supporting belief; i.e., it can only be justified by its relation to still another justified belief; and so on ad infinitum. The standard coherentist response to this difficulty is to opt for a circle of justification, rather than an infinite regress, and then to switch from local to holistic justification. I find this response quite inadequate, but this is not the place to go into that. Instead I want to stress the difference in the difficulty entailed by the higher level regress. The preference for a circle over an infinite set is of no avail here. Since there is a regress of *levels* we are foreclosed from doubling back. No adequate support-belief at an earlier stage will serve to do the job required at a later stage because it will have the wrong content. At each stage what is required is a justified belief to the effect that the "reason for" relationship at the *immediately previous stage* is an adequate one; and no earlier beliefs of that sort in the hierarchy will have been concerned with that particular

"reason for" relationship. Hence there is no alternative here to an infinite regress. And, needless to say, it is highly doubtful that any of us is in possession of such an infinite hierarchy of "adequate support" beliefs.

III

PI has not emerged in strong shape from our examination. Let's turn now to the second construal of an internalist constraint on justification, and see if it fares any better. This second construal has to do with the kind of access we can have to justifiers. The general idea is that possible justifiers are restricted to items to which we have a specially favored access. This special access is variously specified as direct, incorrigible, and obtainable just by reflecting. We have already seen Goldman, in "The Internalist Conception of Justification", identifying internalism with PI. Here is an AI formulation from the same essay.

The basic idea of internalism is that there should be guaranteed epistemic access to the correctness of a DDP. No condition of DDP-rightness is acceptable unless we have epistemic access to the DDP that in fact satisfies the condition, i.e., unless we can tell which DDP satisfies it. The internalist's objection to externalism's condition of rightness, i.e., actual optimality, is precisely that cognizers may have no way of telling which DDP satisfies it. Internalism's *own* condition of rightness must, therefore, be such that any cognizer *can tell* which DDP satisfies it.³¹

Another person we cited as a source of PI, Kent Bach, also brings AI into the picture in "A Rationale for Reliabilism".

Internalism . . . treats justifiedness as a purely internal matter: if *p* is justified for *S*, then *S* must be aware (or at least be immediately capable of being aware) of what makes it justified and why.³²

The parenthetical phrase formulates the AI alternative to the PI formulation outside the parentheses.

However I have found the most elaborate developments of this conception in epistemologists who do not actually employ the "internalism" label. Thus R. M. Chisholm, in a well known passage, lays it down that whenever we are justified in a belief we can determine by reflection what it is that so justifies us.

We presuppose, second, that the things we know are justified for us in the following sense: *we* can know what it is, on any occasion, that constitutes our grounds, or reason, or evidence for thinking that we know.

In beginning with what we think we know to be true, or with what, after reflection, we would be willing to count as being evident, we are assuming that the truth we are seeking is "already implicit in the mind which seeks it, and needs only to be elicited and brought to clear reflection".³³

Carl Ginet gives a more elaborate statement of this version of internalism

Every one of every set of facts about S's position that minimally suffices to make S, at a given time, justified in being confident that p must be *directly recognizable* to S at that time. By 'directly recognizable' I mean this: if a certain fact obtains, then it is directly recognizable to S at a given time if and only if, provided that S at that time has the concept of that sort of fact, S needs at that time only to reflect clear-headedly on the question of whether or not that fact obtains in order to know that it does.³⁴

In the interest of securing a definite target let's focus on the version of special access internalism that requires *direct* access for justifiers, construed along Ginet's lines. I shall refer to this second construal of internalism as "access internalism" (hereinafter "AI").

Our next order of business should be to consider the relation between the two internalisms. Now that we have completed the laborious process of explicating and refining our conception of PI we are at last in a position to do this. Are the two conceptions importantly different? Just how are they related? Can one be subsumed under the other? Does one imply the other?

First let's consider the possibility that PI is a special case of AI. Is the restriction of justifiers to the subject's viewpoint a special case of a restriction of justifiers to what is directly accessible? Only if one's own perspective is directly accessible, and this does not seem to be the case. The sum total of my justified beliefs cannot be depended on to spread themselves before my eyes on demand, not even that segment thereof that is relevant to a particular belief under consideration. I may know something that provides crucial evidence for p and yet fail to realize this even on careful reflection. We need not invoke Freudian blockages to illustrate this, though they are relevant. It may be that the sheer volume of what I know about, e.g., ancient Greek philosophy, is too great for my powers of ready retrieval; or some of this material may be so deeply buried as to require special trains of association to dislodge it. We are

all familiar with cases in which something we knew all along failed to put in an appearance when it was needed to advance a particular inquiry. And, remembering our last modification of PI, still less is it the case that *what I would be justified in believing had I been behaving as I ought* is readily available on reflection.

Thus an item may pass the PI test without passing the AI test. PI is not a special case of AI. How about the converse? Is the restriction to the directly accessible just a special case of the restriction to the subject's justified beliefs and knowledge? Only if nothing other than my knowledge and justified beliefs is directly accessible to me. But that is clearly not the case. My feelings and other conscious experiences are directly accessible if anything is. And even if it were true, as I see no reason to suppose it to be, that I cannot have a conscious experience without knowing that I do, still the experience is distinguishable from the knowledge of the experience. Hence an item can pass the AI test without passing the PI test. This is what makes it possible for partisans of AI like Chisholm and Ginet to recognize immediate justification and to escape coherentism.

Thus PI and AI look quite independent of one another. But surely they must be closely related in some way. Otherwise how can we understand the fact that they are so persistently lumped together under the "internalism" label? And in fact on closer inspection we can see an interesting connection. We can think of AI as a broadening of PI. Whereas PI restricts justifiers to what the subject already justifiably believes (or, in the modified version, to that plus some of what the subject would justifiably believe under ideal conditions), AI enlarges that to include what the subject *can* come to know just on reflection. It is clear that any item that passes the AI test is something that is readily assimilable into the subject's viewpoint, just on reflection. AI, we might say, enlarges the conception of the subject's perspective to include not only what does in fact occur in that perspective (and what should occur), but also what *could* be there if the subject were to turn his attention to it.

Next let's turn to what can be said in support of AI. We have seen that PI is most plausibly supported on a deontological conception of justification, and the AI constraint has also been defended on the same conception. Here we are fortunate to have an explicit statement of the argument from Carl Ginet.

Assuming that S has the concept of justification for being confident that p, S *ought* always to possess or lack confidence that p according to whether or not he has such justification. At least he ought always to withhold confidence unless he has justification. This is simply what is

meant by having or lacking *justification*. But if this is what S ought to do in any possible circumstance, then it is what S *can* do in any possible circumstance. That is, assuming that he has the relevant concepts, S can always tell whether or not he has justification for being confident that p. But this would not be so unless the difference between having such justification and not having it were always directly recognizable to S. And that would not be so if any fact contributing to a set that minimally constitutes S's having such justification were not either directly recognizable to S or entailed by something directly recognizable to S (so that its absence would have to make a directly recognizable difference). For suppose it were otherwise: suppose that some part of a condition minimally sufficient for S's being justified in being confident that p were *not* entailed by anything directly recognizable to S. Then S's position could change from having such justification to lacking it without there being any change at all in what is directly recognizable to S. But if there is no change in directly recognizable features of S's position, S cannot tell that his position has changed in other respects: no matter how clearheadedly and attentively he considers his position he will detect no change. If it seemed to S before that he had justification for being confident that p then it must still seem so to him. So this sort of justification would be such that it would not always be possible for its subject to tell whether or not he possessed it, which is contrary to what we noted is an obvious essential feature of justification. So there can be no such justification. That is, there can be no set of facts giving S justification for being confident that p that has an essential part that is neither directly recognizable to S nor entailed by something directly recognizable to S.³⁵

Note that the conclusion of this argument is not quite the same as the AI thesis I previously quoted from Ginet. According to that thesis, every part of a justifier must be directly recognizable; but the argument purports to show only that a justifier must be either this or *entailed* by what is directly recognizable. Ginet may feel that the additional disjunct makes no significant difference, but this is not the case. One may not be able to spot everything that is entailed by what is directly recognizable; the disjunctive conclusion leaves open the possibility of justifiers that are not wholly identifiable from what is directly recognizable. However I shall suppress this difficulty in the ensuing discussion. For the sake of simplicity I shall consider the thesis in the simpler form, bringing in the second disjunct only where it is specially relevant to the point under consideration.

I have said that Ginet argues from a deontological conception of justification, but this may not be obvious from his formulation of the argument. I shall try to make it more obvious. But first let's note that Ginet explicitly lays out such a conception.

One is *justified* in being confident that p if and only if it is not the case that one ought not to be confident that p: one could not be justly reproached for being confident that p.³⁶

This concept does not explicitly appear in the argument, but it is just below the surface. Ginet uses this concept to define the concept of *having a justification* that he employs in the argument.

I shall take "S has justification for being confident that p" . . . to mean S is in a position such that if he is, or were to be, confident that p then he is, or would be, justified in being so.³⁷

We then get "is justified in being confident that p" defined deontologically as in the previous quotation. Thus the concept used in the argument is, so to say, the first derivative of a deontological conception. It is the concept of having what it takes to be justified in the deontological sense if one will only make use of those resources.

Before entering into a critical scrutiny of the argument, let's note some of its features, with special attention to the points we were making concerning the argument for PI. First, the argument should, by rights, apply to overrides of *prima facie* justification as well as to justifiers. Consider that done. Second, Ginet is obviously presupposing direct voluntary control of belief. Since "in any possible circumstance", "S *ought* always to possess or lack confidence that p according to whether or not he has such justification", this is something that "S *can* do in any possible circumstance". It is always possible for S to stop and consider any actual belief of his, or any candidate for belief, and bring it about then and there that he does or does not adopt or continue to hold the belief according as he has or lacks sufficient justification for it.³⁸ It is not so clear whether Ginet's concept of justification applies only to beliefs that are acquired by a deliberate choice, and then only in terms of what is true at the moment of acquisition. Let's suppose that he is only assuming the ever present possibility of a deliberate choice between adopting (continuing) a belief and refraining from doing so, and that to be justified in believing that p is to be so situated that if one were, in that situation, to choose to believe that p (or continue to do so) one could not be blamed, on intellectual grounds, for that choice.

It will help us to critically evaluate Ginet's argument if we exhibit its skeleton.

- 1 S ought to withhold belief that p if s/he lacks justification for p.³⁹
- 2 What S ought to do S can do.
- 3 Therefore, S can withhold belief wherever S lacks justification.
- 4 S has this capacity only if S can tell, with respect to any proposed belief, whether or not S has justification for it.
- 5 S can always tell us this only if justification is always directly recognizable.
- 6 Therefore justification is always directly recognizable.

This bare bones rendition should make it apparent where the argument goes astray. It is at step 5. Step 5 claims that S can tell whether he has justification for a belief only if it is directly recognizable by him whether he does or not. But why should we suppose this? Ginet, in company with almost all contemporary epistemologists, wisely avoids holding that one can know only what is evident to one on simple reflection and what is entailed by that. We know many things only because we have reasons for them in the shape of other things we know, and these reasons are not always deductively related to what they support. Thus direct recognition is only one way to acquire knowledge. Why should we suppose that only this way is available for knowing about justification? That would have to be argued. In the absence of any such argument we are at liberty to deny that justification can always be spotted just by reflection. The argument leaves standing the possibility that S might, in various instances, come to know in some other way whether he has a justification for p.

Consider the ethical analogy that is inevitably suggested by Ginet's argument. There is an exactly parallel argument for the thesis that the justification of actions is always directly recognizable. But that is clearly false. Often I have to engage in considerable research to determine whether a proposed action is justified. If it is a question of whether I would be justified in making a certain decision as department chairman without consulting the executive committee or the department as a whole, I cannot ascertain this just by reflection, unless I have thoroughly internalized the relevant rules, regulations, by-laws, and so on. Most likely I will have to do some research. Would I be legally justified in deducting the cost of a computer on my income tax return? I had better look up the IRS regulations and not just engage in careful reflection. The situation is similar with respect to more strictly moral justification. Would I be morally justified in resigning my professorship as late as April 12 in order to accept a position elsewhere for the following fall? This depends, *inter alia*, on how much inconvenience this would cause my present

department, what faculty resources there are already on hand for taking up the slack, how likely it is that a suitable temporary replacement could be secured for the coming fall; and so on. There is no guarantee that all these matters are available to me just on simple reflection. Why should we suppose, without being given reasons to do so, that the justification of beliefs is different in this respect?

Let's remember that in the argument we quoted Ginet supported his position by a *reductio* that runs as follows.

- 1 Suppose that some part of a justification were not entailed by what is directly recognizable to S.
- 2 Then S's position could change from having such justification to lacking it without there being any change in what is directly recognizable to s.
- 3 But then S cannot tell that his position *vis-à-vis* justification has changed.
- 4 Therefore if S can always tell what his justificatory situation is, no part of a justification can fail to be directly recognizable.

This argument, in step 3, presupposes a strong foundationalism according to which any knowledge I can have is based on what is directly recognizable to me, and this could well be contested. But even if we go along with this, the argument is unsound. The trouble is in 2, in the assumption that *anything* not entailed by the directly recognizable can change with no change in what is directly recognizable. To assume this is to assume that the non-directly recognizable is effectively reflected in what is directly recognizable only if the former is entailed by the latter. For if there are other modes of reflection, then a change in the former will sometimes be mirrored in a change in the latter, even when the former is not *entailed* by the latter. For convenience of exposition, let's lump together everything that is not entailed by anything directly recognizable by me as "the world." It is certainly the better part of reason to recognize that much of the world is not adequately reflected in what I can directly recognize; if that were not the case I would be in an immeasurably stronger epistemic position than is the lot of humanity. But to suppose that the world beyond my direct recognition *never* reveals itself in what I can directly recognize would be subversive of the very type of foundationalism this argument presupposes. For in that case the foundations would ground no knowledge of anything beyond themselves except by way of logical deduction. And I am sure that Ginet does not want that. If then a change in "the world" is sometimes reflected in changes in the directly recognizable, why suppose that this is not the case with respect to justification?

Put the matter another way. All that Ginet can extract from his strong foundationalist assumption, his deontological concept of justification, and the

“ought implies can” principle, is that it is always possible to determine *from* what is directly recognizable to the subject whether the subject is justified in a certain belief. But that does *not* imply that what does the justifying is itself directly recognizable, or is entailed by what is directly recognizable. It only implies that either it has this status *or* it can be ascertained on the basis of what is directly recognizable.⁴⁰

However, Ginet’s argument can easily be transformed into an argument for a more moderate form of AI. To begin with the other extreme, suppose we formulate AI just as the view that to be a justifier an item must be epistemically accessible in some way to the subject. It is not *impossible* for the subject to acquire that bit of knowledge (or justified belief). It does seem that Ginet’s argument would establish that much accessibility, granted his premises. If I ought to do something that requires knowing the answer to a certain question, it must be *possible* for me to get that answer.

But what is the significance of this result? What does this constraint exclude? It excludes factors that are in principle unknowable by human beings; but it is dubious that any of the parties to the discussion are disposed to suggest justifiers that satisfy that description. The putative justifiers that internalists typically wish to exclude are items other than beliefs and experiences of the subject. Bonjour’s clairvoyant subject in “Externalist Theories of Empirical Knowledge” is representative of the disputed territory. This person in fact has clairvoyant powers but has neither any understanding of what is going on nor any good reasons for supposing that these powers are reliable. So far as he can tell, the beliefs simply occur to him, and he is, strangely enough, irresistibly constrained to accept them. What shall we fasten on as the strongest candidate for a justifier here? There are no beliefs or experiences on which the clairvoyant beliefs are based. Let’s say that if anything justifies them it is their resulting from the exercise of reliable clairvoyant powers. The subject knows nothing of such powers. But is it *impossible* that he should discover them and discover that they are reliable? I see no reason to suppose that. He might ascertain this just by discovering that these strange beliefs about distant places that apparently just popped into his mind out of nowhere are invariably true. It appears, then, that the requirement of being knowable somehow is too weak to be of much interest.

Perhaps there is a mean between the extremes that is both of some significance and still not too strong to be supportable. We might try requiring knowability, not just on reflection at the moment, but at least without a great deal of research. Admittedly this is quite vague. The vagueness may be reduced by bringing in the notion of what could reasonably be expected in the way of time and effort devoted to searching out the justifiers. These expect-

tations might differ from case to case, depending on the kind of justifiers that would be required, the capacities and initial position of the subjects, and so on. If a belief is based on experience, we would naturally expect the subject to ascertain that right off the bat. If, on the other hand, a belief is based on a large and complex body of evidence, we would not expect the subject to be able to survey all that in a moment. And so on. We might dub this intermediate conception “reasonably immediate accessibility.”⁴¹ Although this may seem a more reasonable requirement than Ginet’s, and although it obviously is less restrictive, this increase in modesty has not purchased any greater support by Ginet’s line of argument. I can’t see that an “ought implies can” principle supports a “reasonably immediate accessibility” any more than it supports a direct recognizability. In the absence of further reasons to the contrary, all that would seem to be required by the principle is knowability in some way or other.

We should also consider whether a Ginet-like argument can be mounted on the basis of an indirect control version of a deontological conception of justification. In a word, NO. Ginet’s argument depends on the point that knowledge is required to carry out intellectual obligations. But on the indirect control version intellectual obligations do not attach to believing or refraining from believing. Hence the knowledge needed to carry out these obligations is not knowledge of whether there is a justification for some particular belief.

Now let’s turn to the question of a higher level extension of AI. It is clear that the AI constraint, like the PI constraint, can be imposed on various levels. We saw that the basic argument for PI equally supported the first and second level constraints. It supported both the claim that a justifier had to be a justified belief, and the claim that one justified belief can justify another only if the subject is justified in the higher level belief that the first belief does adequately support the second. What about the argument for AI? Ginet does not use his argument to support a higher level extension. As noted earlier, he does impose a higher level PI constraint on mediate justification, but he associates no higher level constraint of any kind with his AI position. He takes AI to require only that *justifiers* be directly recognizable, not that it be directly recognizable that they possess justificatory efficacy. And yet his argument supports a higher level AI requirement just as strongly, or weakly, as the lower level requirement. This can be seen as follows. Suppose that the sorts of things that can count as justifiers are always accessible to me, but that it is not always accessible to me which items of these sorts count as justifications for which beliefs. I have access to the justifiers but not to their justificatory efficacy. This will take away my ability to do what I am said to have an obligation to do

just as surely as the lack of access to the justifiers themselves. To illustrate, let's suppose that experiences can function as justifiers, and that they are accessible to us. I can always tell what sensory experiences I am having at a given moment. Even so, if I am unable to tell what belief about the current physical environment is justified by a given sensory experience, I am thereby unable to regulate my perceptual beliefs according as they possess or lack experiential justification. Knowing what the facts are doesn't suffice for enabling me to regulate my behavior accordingly; I also have to know the significance of these facts for what I ought to do. Thus the "ought implies can" argument supports the higher level requirement to just the extent to which it supports the lower level requirement.

Thus AI, too, has higher-level troubles. The trouble is not nearly as severe as its PI analogue. For one thing, what is required here is not actual higher level knowledge (justified belief) about justification, but only the capacity to obtain it. Thus we are not required to attribute to all subjects an absurdly inflated body of actual knowledge about the conditions of justification. Second, for the same reason we are not faced with nasty infinite regresses or hierarchies. Since to be justified in believing that *p*, *S* need not actually justifiably believe that the alleged justifier is fitted to do its job, but only be capable of ascertaining this, we are not committed to an actual infinite hierarchy of such justified beliefs. Nevertheless there are serious questions as to whether even a modest AI higher level requirement is not too severe. The requirement implies that state of affairs, *A*, cannot justify me in believing that *p* unless I am capable of determining that *A* is a genuine justification for a belief that *p*. But how many subjects are capable of this? Indeed, there are substantial grounds for scepticism about the possibility of anyone's having adequate reasons for claims about justification. The grounds I have in mind concern the specter of epistemic circularity, the danger that, e.g., any otherwise promising argument for a principle laying down conditions under which perceptual beliefs are justified will have to use perceptual beliefs among its premises. I have considered this problem elsewhere and have concluded that, despite the pervasive presence of epistemic circularity in such arguments, it is possible to be justified in beliefs about the conditions of justification.⁴² But even if that rather optimistic conclusion is justified, it still seems that many subjects are not capable of acquiring adequately justified beliefs concerning what justifies what. To go into this properly we would have to decide what it takes for the justification of such beliefs, and there is no time for this lengthy investigation in this paper. Let me just say that it seems eminently plausible that beliefs about what justifies what would have be justified by reasons (not directly justified), and it would seem that such reasons are directly accessible to few, if any of us.

All this suggests limiting AI to the lower level. Something can function as a justifier only if it is (fairly readily) accessible, but in order to function as a justifier it is not necessary that its justificatory efficacy be likewise accessible. At some point we must rely on things just *being* a certain way, without its also being the case that we do or can assure ourselves that they are that way. And this would seem to be the proper place to draw that line. We shall return to this possibility in the last section. For now, let's sharpen the issue by recalling the fact that a reliability account of justification (*S* is justified in believing that *p* if and only if *S*'s belief that *p* was reliably produced) is often attacked on the grounds that justification could not be lost by a loss of reliability, so long as the situation is the same, *so far as we can tell*. Consider a possible world that is indistinguishable from the actual world so far as we can tell, but in which a Cartesian demon has rigged things so that our perceptual beliefs concerning external physical objects are all false, since there are no such objects. Since such a world is indistinguishable (by us) from our world, we would have just as much justification for our perceptual beliefs there as we actually do. But *ex hypothesi* those beliefs would not be reliably formed. Hence reliability is not necessary for justification. Here are some snatches of such an argument from an essay by Richard Foley. (The demon world is called "w".)

If we are willing to grant that in our world some of the propositions *S* perceptually believes are epistemically rational, then these same propositions would be epistemically rational for *S* in *w* as well. After all, world *w* by hypothesis is one which from *S*'s viewpoint is indistinguishable from this world. So, if given *S*'s situation in this world his perceptual belief *p* is rational, his belief *p* would be rational in *w* as well.

Even if, contrary to what we believe, our world is world *w*, it still can be epistemically rational for us to believe many of the propositions we do, since the epistemic situation in world *w* is indistinguishable from the epistemic situation in a world which has the characteristics we take our world to have. The point here is a simple one. In effect, I am asking you: aren't some of the propositions you believe epistemically rational for you to believe? And wouldn't whatever it is that make those propositions epistemically rational for you also be present in a world where these propositions are regularly false, but where a demon hid this from you by making the world from your viewpoint indistinguishable from this world (so that what you believed, and what you would believe on reflection, and what you seemed to remember, and what you experienced were identical to this world)?⁴³

In each of these passages the fact that we cannot distinguish *w* from the actual world is taken to imply that whatever justifies a certain belief in the one world will ipso facto justify that same belief in the other world. This argument presupposes an AI internalist constraint on both levels. For suppose AI put constraints only on what can count as a justifier, not also on what has justificatory efficacy for which beliefs. In that case the reliabilist would remain free to claim that although the same putative justifiers (of perceptual beliefs) are present in the two worlds, they do justify perceptual beliefs in the actual world but not in *w*, since their production of perceptual beliefs is reliable in the actual world but not in *w*. If and only if justificatory efficacy were subject to an AI constraint would this be impossible, as Foley claims. If, on the other hand, one follows my suggestion that we adopt an accessibility constraint only on the lower level, we can recognize that a state of affairs, *A*, can justify a belief that *p* in one possible world and not in another, even though we can't tell any difference between the two worlds.

IV

The upshot of the paper is that existing forms of internalism are in serious trouble. Both PI and AI run into severe difficulties over their higher level component, but if we try shearing off that component we lose such support as has been provided them. That support is less than impressive in any case. The only arguments of any substance that have been advanced proceed from a deontological conception of justification and inherit any disabilities that attach to that conception. Indeed, PI gains significant support only from the most restrictive form of a direct voluntary control version of that conception, one that is, at best, of limited application to our beliefs. As for AI, the arguments in the literature that are designed to establish a direct recognizability version markedly fail to do so. And it is not clear that a more moderate form of AI can be developed that will be both well supported by these arguments and strong enough to have any cutting edge.

Thus internalism has not emerged in strong shape from this examination. It looks as if no sort of internalist constraint can be justified, and hence that an unrestricted externalism wins the day. However I do not believe that so extreme a conclusion is warranted. I am convinced that the considerations advanced in this paper show that existing versions of internalism are untenable, and that such arguments as have been advanced for them fail to establish any form of that position. And yet I am inclined to suppose that a suitably modest form of AI internalism can be supported, though in a very different

way from any employed by the internalists we have been discussing. If any readers have persevered this far, I will not further test their patience by embarking on a full dress development and defense of this suggestion, but I will just indicate what I have in mind.

Earlier I indicated that what I called a strong position (SP) conception of justification does not support any sort of internalist restriction. One can believe that *p* in such a way as to be in a strong position to acquire the truth whether or not that belief is supported adequately by other of one's justified beliefs (PI), and whether or not one has strong epistemic access to the grounds for the belief. In my "Concepts of Epistemic Justification" I have argued for the superiority of the PS conception over any kind of deontological conception. Thus, so far as these options for a concept of justification are concerned, pure externalism reigns supreme. Nevertheless I do not take this to be the last word. Even if internalist intuitions cannot be supported by the most basic features of the concept of justification, they may have a certain validity on their own, as an independent contribution to the concept. Let's once more consider "out of the blue" reliable modes of belief formation. Let's say that when I am suddenly seized with apparently irrational convictions concerning the current weather in some distant spot, these convictions always turn out to be correct. If there is nothing to justification other than believing in such a way as to be in a strong position to acquire the truth, then we should say that I am justified in those convictions. And yet we are loath to admit this, at least before I become aware of the reliability of this mode of belief formation. (After I become aware of this I have an adequate reason for the convictions, and this should satisfy any internalist scruples.) Why this reluctance? What is missing? What is missing, of course, is any basis or ground that *S* has, possesses, for his belief, anything that he can point to or specify as that which gives him *something to go on* in believing this, any *sign* or *indication* he has that the belief is true. Wherever nothing like this is involved, we feel uneasy in taking *S*'s belief to be *justified*. Thus it looks as if there is a basic, irreducible, requirement of *epistemic accessibility of ground for the belief* that attaches to our concept of epistemic justification.⁴⁴ For reasons we have rehearsed at some length, let's take the accessibility required to be of the relatively modest sort that we earlier called "reasonably immediate accessibility".

Can this requirement be derived from other features of the concept? It certainly cannot be derived from an SP conception, and we have seen that such support as it gleans from a deontological conception would bring fatal difficulties with it, even if such a conception were viable for epistemology. I am inclined to think that the requirement is a fundamental constituent of our concept of epistemic justification, though I do not take that to imply that

there can be no sort of explanation for its presence. I will conclude by briefly adumbrating what I take to be responsible for this internalist feature of the concept.

My suggestion is that the background against which the concept of epistemic justification has developed is the practice of critical reflection on our beliefs, the practice of the epistemic assessment of beliefs (with respect to the likelihood of their being true), the challenging of beliefs and responses to such challenges. To respond successfully to such a challenge one must specify an adequate ground of the belief, a ground that provides a sufficient indication of the truth of the belief. It would, of course, be absurd to suggest that in order to be epistemically respectable, laudatory, or acceptable (justified) a belief must have actually been put to such a test and have emerged victorious. In suggesting that the concept has developed against the background of such a practice the idea is rather that what it is for a belief to be justified is that the belief and its ground be such that it is in a position to pass such a test; that the subject has what it takes to respond successfully to such a challenge.⁴⁵ A justified belief is one that *could* survive a critical reflection. But then the justifier must be accessible to the subject. Otherwise the subject would be in no position to cite it as what provides a sufficient indication that the belief is true. This, baldly stated, is what I take to be the explanation of the presence of an AI internalist constraint in the concept of epistemic justification. Further development of this suggestion must await another occasion.

Notes

- 1 See e.g., Laurence Bonjour, "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy*, Vol. V (1980).
- 2 Bonjour, "Externalist Theories", p. 55.
- 3 p. 55.
- 4 p. 56.
- 5 *The Monist*, Vol. 68, no. 2 (April, 1985), p. 247.
- 6 *Midwest Studies in Philosophy*, Vol. V (1980), p. 32. Later in this essay Goldman considers what conditions should be laid down for the acceptance of a *doxastic decision principle* (DDP). A DDP is a "function whose *inputs* are certain conditions of a cognizer – e.g., his beliefs, perceptual field, and ostensive memories – and whose *outputs* are prescriptions to adopt (or retain) this or that doxastic attitude . . ." (p. 29) Here is what he takes to be "the condition appropriate to externalism":

(1) DDP X is right if and only if: X is *actually* optimal.

Whereas the first shot at formulating an appropriate condition for internalism is the following:

(2) DDP X is right if and only if: we are *justified* in believing that X is optimal. (pp. 33–4)

- 7 Note that in the passage quoted above Kent Bach says that "internalism requires that a person have 'cognitive grasp' of whatever makes his belief justified". This too would seem to allow that what makes the belief justified could be an item of any (suitable) sort, provided the person has a "cognitive grasp" of it.
- 8 Thus although in that passage Bonjour is arguing against the existence of "basic beliefs", i.e., immediately justified beliefs, the argument, if successful, will rule out only *purely* basic beliefs. It will not rule out beliefs a part of whose justification consists in something other than justified beliefs of the same subject.
- 9 What I am calling "Bonjour's version" does not represent his considered position, which is more like the other version. However the former is suggested by the passage under discussion. Perhaps Bonjour was led into it there because he was arguing with a partisan of immediate knowledge who claims that a certain non-belief is sufficient for the justification of a certain belief. Having no reason to deny that the non-doxastic state of affairs obtains, Bonjour simply confined himself to alleging that even if it does obtain, the subject will also have to be justified in believing that it obtains.
- 10 I shall use "overrider" for something that cancels out a *prima facie* justification. Unlike some theorists I shall refrain from using "defeater" for this purpose, saving that term (though not using it in this paper) for a fact the mere holding of which prevents a true, overall justified belief from counting as knowledge.
- 11 Thus principle (B) in R. M. Chisholm, *Theory of Knowledge*, 2nd edn. (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1977) runs as follows.

(B) For any subject S, if S believes, without ground for doubt, that he is perceiving something to be F, then it is beyond reasonable doubt for S that he perceives something to be F.

And "ground for doubt" is explained as follows.

(D4.3) S believes, *without ground for doubt*, that p = Df (i) S believes that p and (ii) no conjunction of propositions that are acceptable for S tends to confirm the negation of the proposition that p. (p. 76)

The PI constraint comes in by requiring "grounds for doubt" that consists in propositions that are "acceptable" for the subject, in order that the *prima facie* justification of perceptual beliefs be overthrown.

- 12 "A Plethora of Epistemological Theories", in George S. Pappas, ed., *Justification and Knowledge* (Dordrecht: D. Reidel Pub. Co., 1979), p. 101.
- 13 To be sure, there are more alternatives than the ones we have mentioned. In his book, *Knowledge* (Oxford: Clarendon Press, 1974), Keith Lehrer pumps for a coherence theory in which the test of justification is coherence, not with the actual set of beliefs of the subject but with what Lehrer calls the subject's "corrected doxastic system", that subset "resulting when every statement is deleted which describes S as believing something he would cease to believe as an impartial and disinterested truth-seeker". (p. 190)
- 14 "Can Empirical Knowledge Have a Foundation?", *American Philosophical Quarterly*, 15, 1 (Jan., 1978), p. 5.
- 15 Of course the "feature" could be so specified that it included the subject's justified beliefs about another feature. But then it would be this latter features with respect to which BonJour is requiring the justified beliefs, and the point would still hold.
- 16 See, e.g., Richard Foley, "What's Wrong with Reliabilism?", *The Monist*, Vol. 68, no. 2 (April, 1985).
- 17 This requirement for mediate justification is embraced by many epistemologists who do not advocate (4) with its denial of any immediate justification. See, e.g., Carl Ginet, *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel Pub. Co., 1975), pp. 47-9.
- 18 BonJour, "Externalist Theories", p. 63.
- 19 Keith Lehrer, op. cit., pp. 187-8.
- 20 Pollock, "A Plethora of Epistemological Theories", p. 106. This does not represent Pollock's overall view.
- 21 This argument may be thought of as a development of BonJour's suggestion that the subject must "possess the justification" for the belief "if his acceptance of it is to be epistemically rational or responsible". ("Externalist Theories . . .", p. 55.)
- 22 Note that what this argument supports is a *positive, local justification* version of PI. But precisely parallel arguments can be given for other versions. For the suggestion of such an argument for a *negative, local justification* version, see the quotation from Wolterstorff in the following footnote. For an argument for a mere belief version, whether local or holistic, see Pollock's "A Plethora of Epistemological Theories", p. 109.
- 23 Here are some adumbrations of this argument. ". . . on the externalist view, a person may be ever so irrational and irresponsible in accepting a belief, when judged in light of his own subjective conception of the situation, and may still turn out to satisfy Armstrong's general criterion of reliability. This belief may in fact be reliable, even though the person has no reason for thinking that it is reliable . . . But such a person seems nonetheless to be thoroughly irresponsible from an epistemic standpoint in accepting such a belief, and hence not justified, contrary to externalism." (BonJour, "Externalist Theories . . .", p. 59.) Here is

- another adumbration, this time from the standpoint of a negative coherence theory that holds a belief to be justified provided one has no sufficient reason for giving it up. "If a person does not have adequate reason to refrain from some belief of his, what could possibly oblige him to give it up? Conversely, if he surrenders some belief of his as soon as he has adequate reason to do so, what more can rightly be demanded of him? Is he not then using the capacities he has for governing his beliefs, with the goal of getting more amply in touch with reality, as well as can rightly be demanded of him?" (Nicholas Wolterstorff, "Can Belief in God Be Rational?", in Alvin Plantinga and Nicholas Wolterstorff, eds., *Faith and Rationality* (Notre Dame, IN: University of Notre Dame Press, 1983), p. 1963. Note the crucial occurrence in these passages of terms like "irresponsible", "oblige", and "rightly demanded". Both these authors, as well as other PI internalists, note the parallel between what is required for epistemic and for ethical justification. In both cases, it is argued, what is required is that the belief or the action be the one to adopt, so far as one can tell from one's own viewpoint on the world.
- 24 "Concepts of Epistemic Justification", *The Monist*, Vol. 68, no. 2 (Jan., 1985). For other developments of the deontological conception see Ginet, *Knowledge, Perception, and Memory*, Wolterstorff, "Can Belief in God Be Rational", and M. B. Naylor, "Epistemic Justification", 1978, unpublished.
- 25 BonJour is an exception in trying to combine features of the two conceptions. On the one hand, he argues for PI from a deontological conception of justification. On the other hand, as we have seen, he presupposes the truth-conductivity of justification in formulating his higher-level requirement.
- 26 Thus Wolterstorff: "Locke assumes - rightly in my judgment - that we have an obligation to govern our assent with the goal in mind of getting more amply in touch with reality." Wolterstorff, "Can Belief in God Be Rational?", p. 145.
- 27 There are other ways in which a subject's epistemic situation might have been different from what it actually is had the subject been doing a better job of carrying out his/her intellectual obligations. In particular, the subject might have had justifiers that s/he does not actually possess. However it is not at all clear that this and other differences from the actual situation have the same bearing on justification as the lack of overrides that one should have had. Consider a case in which if I had been attending to the matter as I should have I would have had justified beliefs that adequately support the belief that Jones is untrustworthy. As things actually stand I do not have adequate reasons for supposing that. Here, going on the justified beliefs I actually have, we would have to say that I would not be justified in believing that Jones is untrustworthy. But nor does it seem that that judgment would be reversed by the consideration that I would have had adequate support had I been conducting myself properly. Surely we don't want to say that the thing for me to do is to adopt that belief *in the absence of sufficient reasons*, whatever would have been the case had I been satisfying my intellectual obligations.

- 28 At a few points in the preceding exposition the need for this qualification was more or less evident. Thus at one point I represented the deontologist as saying that the justification of a given belief depends on the "best representation of the world of which I am currently capable". I have also used such phrases as "one's best judgment of the facts" and "so far as one can tell". All of these phrases point to the "ideal viewpoint" rather than to the actual viewpoint. The best representation of the world of which I am currently capable may not be the representation I actually have. There will be a discrepancy provided, as is usually the case to some extent, I have not made full use of my opportunities for ascertaining relevant features of the world. The importance of overrides that a subject ought to have but doesn't is well brought out by Wolterstorff, "Can Belief in God Be Rational?", pp. 165-6.
- 29 For a discussion of this issue see my "Concepts of Epistemic Justification".
- 30 If this latter sort of reason is required, that constitutes a serious stumbling block, for it seems that we are rarely justified in any such belief, unless the requirements for justification are set very low. But that is not our present concern.
- 31 Goldman, "The Internalist Conception of Justification", p. 35. Remember that a DDP is, roughly, a principle that declares certain beliefs to be justified under certain conditions. Therefore the requirement that there be maximal epistemic access to a DDP is an AI analogue of what we were calling the "higher level requirement" component of PI. Interestingly enough, when it comes to a high accessibility "lower level requirement" with respect to justifiers, "input to the DDP" in Goldman's lingo, Goldman lays this down on his own, with no hint that it is required by internalism as contrasted with externalism. "If a DDP is to be actually *usable* for making deliberate decisions the conditions that serve as inputs must be *accessible* or *available* to the decision-maker at the time of decision. The agent must be *able to tell*, with respect to any possible input condition, whether that condition holds at the time in question." (p. 30) He even spells this out in such a way that it is *infallible* access that is required. "But what exactly do we mean in saying that a person 'can tell' with respect to a given condition whether or not that condition obtains? Here is a reasonable answer: 'For any person S and time t, if S asks himself at t whether condition C obtains at the time in question, then S will believe that condition C obtains then if and only if it does obtain then.'" (31).
- 32 K. Bach, "Rationale for Reliabilism", p. 250. Cf. p. 252.
- 33 Chisholm, *Theory of Knowledge*, p. 17. The quotation is from C. I. Lewis, *Mind and the World Order*. It should be acknowledged that in a later essay Chisholm states this assumption only for "some of the things I am justified in believing". See "A Version of Foundationalism", *Midwest Studies in Philosophy*, Vol. V (1980), p. 546.
- 34 *Knowledge, Perception, and Memory* (Dordrecht: D. Reidel Pub. Co., 1975), p. 34.
- 35 Ginet, *Knowledge, Perception, and Memory*, p. 36.
- 36 Ginet, *Knowledge, Perception, and Memory*, p. 28.
- 37 Ginet, *Knowledge, Perception, and Memory*, p. 28.

- 38 In "Contra Reliabilism", *The Monist*, Vol. 68, no. 2 (April, 1985), Ginet defends this assumption against objections from myself. Note that Ginet's argument could easily be recast in an "indirect voluntary control" form. Instead of premising that it is always possible to decide whether or not to believe, or to continue believing, that p in the light of the presence or absence of a sufficient justification, one can hold instead that it is always possible to decide whether to do various things to encourage or discourage belief that p, in the light of the presence or absence of a sufficient justification. The direct recognizability of justifiers will be as strongly supported by this version as by the original version.
- 39 Ginet recognizes that we are intellectually obligated to refrain from believing that p in the absence of justification, but he wisely holds back from claiming that we are obligated to believe that p wherever we have a justification. The presence of justification gives me a *right* to believe, but I am not obliged to exercise that right; I have a choice as to whether or not to do so. It seems plausible to hold, e.g., that I am justified in believing everything that is entailed by my justified beliefs. But an infinite set of beliefs is so entailed. Thus if I were obligated to believe everything for which I have a justification I would be in a pretty pickle. Ginet's recognition of this point is evinced by his modifying "S *ought* always to possess or lack confidence that p according to whether or not he has such justification" to "At least he ought always to withhold confidence unless he has justification".
- 40 We could also attack the direct accessibility form of AI by pointing to the fact that not all commonly recognized justifiers satisfy the constraint. Remember that when we were considering the relations of PI and AI we pointed out that one cannot, in general, retrieve all relevant justified beliefs of oneself just on reflection.
- 41 Note that all these accessibility requirements, of whatever degree of stringency, can be thought of as related to PI in the same way. Any item that is epistemically accessible to S can be thought of as potentially an item in S's perspective on the world. Hence any sort of AI can be thought of as a broadening of PI to include potential additions to the perspective, as well as its present constituents.
- 42 See my "Epistemic Circularity", forthcoming in *Philosophy and Phenomenological Research*.
- 43 Richard Foley, "What's Wrong with Reliabilism?", *The Monist*, Vol. 68, no. 2 (April, 1985). See also Carl Ginet, "Contra Reliabilism", *ibid*.
- 44 Since I do not find any like tendency to withhold the concept of justification when the justificatory efficacy of the ground is not readily accessible to the subject, I am not saddled with the burden of a higher level accessibility constraint.
- 45 One indication that this is the right way to think about justification is the fact that we find it incongruous to apply the concept to beings that are incapable of critical reflection on their beliefs. The question of whether a dog is *justified* in supposing that his master is at the door is one that does not seem to arise. There are, to be sure, problems as to just how this restriction is to be interpreted. It

seems clearly all right to apply the concept to human beings that have little skill at the game of challenge and response. The applicability to small children is less clear, but note that in both these cases we are dealing with beings that belong to a species many members of which are capable of critical reflection a full blooded form.

4

HOW INTERNAL CAN YOU GET?

Hilary Kornblith

There has been a debate of late among epistemologists concerning the proper form of a theory of justification. The debate is between those who see themselves as internalists about justification and those who see themselves as externalists about justification. Precisely what it means to be either an internalist or an externalist is itself a matter of debate. In this kind of situation, where the very content of a position becomes an issue, one does well to go back to motivations; one needs to ask why it is that anyone ever defended a version of the position at issue. Only then can one understand what the issues really are.

My focus in this paper will be on internalism. I want to ask, among other questions, why anyone ever wanted to be an internalist and what it is that internalism is supposed to achieve. But one cannot, I think, approach these questions directly. One must first have a specimen internalist on the table, so to speak. I will thus begin by examining some features of Laurence Bonjour's defense of internalism in *The Structure of Empirical Knowledge*. I choose Bonjour's book as a starting point because among its many virtues are its pellucid style and its well-defined motivation. One could not ask for a better entree into the issues which separate internalists from externalists.

In spite of my great admiration for Bonjour's book, I find myself in agreement with very little of it. I will isolate one central point of disagreement, surrounding what Bonjour calls the Doxastic Presumption. By examining the role which this Presumption plays in Bonjour's book, I hope to expose some serious difficulties for Bonjour's form of internalism and, at the same time, to cast some light on the limitations of any possible internalist position.