

JOHN GRECO

INTERNALISM AND EPISTEMICALLY RESPONSIBLE BELIEF

ABSTRACT. In section one the deontological (or responsibilist) conception of justification is discussed and explained. In section two, arguments are put forward in order to derive the most plausible version of perspectival internalism, or the position that epistemic justification is a function of factors internal to the believer's cognitive perspective. The two most common considerations put forward in favor of perspectival internalism are discussed. These are the responsibilist conception of justification, and the intuition that two believers with like beliefs and experiences are equally justified in their like beliefs. In section three it is argued that perspectival internalism is false, and that in fact the position is not supported by a responsibilist conception of justification. Section four explicates two other forms of internalism, which are rejected for reasons similar to those presented against perspectival internalism. In section five, an internalist theory of justification is defended which is supported by a responsibilist conception of justification. Roughly stated, the position is that justified belief is belief which arises from the use of correct rules of reasoning. The idea of correctness is explicated, and the position is distinguished from others which are similar to it.

Contemporary epistemology is involved in a debate about the nature of epistemic justification. On one side of the debate are philosophers who embrace internalist theories of justification. These philosophers agree that, in some important sense, the factors which give rise to justified belief must be internal to the believing subject. On the other side of the debate are philosophers who embrace externalist theories of justification. These philosophers deny that any such privileged relation must exist between the believer and that which is justifiably believed.

Although there are many different versions of internalism, almost all internalists agree that their general position is supported by a "responsibilist" conception of epistemic justification, or a conception of justification which takes epistemic responsibility to be central to justified belief. Internalists argue that epistemically justified belief just is epistemically responsible belief, but since epistemic responsibility is a function of considerations internal to the believing subject, epistemic justification is also a function of such internal considerations. In this paper

Synthese 85, 245–277, 1990.

© 1990 Kluwer Academic Publishers. Printed in the Netherlands.

I will examine several internalist theories, and I will conclude (a) that all of them are false, and (b) that none of them are supported by a responsibilist conception of justification. I will then offer a different theory of epistemic justification, which also may be properly labeled “internalist”, and which is supported by a responsibilist conception. Before turning to considerations of different internalist theories, it will be helpful to say some more about the responsibilist conception of epistemic justification.

1. DEONTOLOGICAL JUSTIFICATION

As noted above, the responsibilist conception takes epistemic responsibility to be central to justified belief. Epistemic responsibility, in turn, is understood in terms of epistemic duty or obligation. Hence this broad conception of justification has also been called “the deontological conception of epistemic justification”.¹ The idea that epistemic justification is centrally concerned with epistemic responsibility has enjoyed a wide acceptance among philosophers. Thus Roderick Chisholm writes,

Let us consider the concept of what might be called an “intellectual requirement”. We may assume that every person is subject to a purely intellectual requirement – that of trying his best to bring it about that, for every proposition *h* that he considers, he accepts *h* if and only if *h* is true. One might say that this is the person’s responsibility *qua* intellectual being One way, then, of reexpressing the locution ‘*p* is more reasonable than *q* for *S* at *t*’ is to say this: *S* is so situated at *t* that his intellectual requirement, his responsibility as an intellectual being, is better fulfilled by *p* than by *q*.²

Laurence Bonjour accepts a similar position in the following passage.

[O]ne’s cognitive endeavors are epistemically justified only if and to the extent that they are aimed at this goal, which means very roughly that one accepts all and only those beliefs which one has good reason to think are true. To accept a belief in the absence of such a reason . . . is, one might say, *epistemically irresponsible*. My contention here is that the idea of avoiding such irresponsibility, of being epistemically responsible in one’s believings, is the core of the notion of epistemic justification.³

As William Alston points out, notions such as epistemic obligation and epistemic responsibility can be understood on the model of their moral analogs. Thus we say that an action of *S*’s is morally justified if and only if *S*’s action does not involve the violation of some moral obligation. Substituting appropriately, we may say that a belief of *S*’s is epistemically justified if and only if *S*’s belief does not involve the violation of

some epistemic obligation.⁴ An example of an epistemic obligation might be to refrain from beliefs not based on good evidence.

As with its moral analog, the concept of epistemic obligation is ambiguous. For we may distinguish between objective and subjective obligations. Thus a parent has an objective moral obligation to provide for the well-being of his child. Whether this obligation is fulfilled depends on whether the child is in fact provided for in an adequate manner. However, the parent also has a subjective obligation to provide for his child. This may be understood in terms of what, from the parent's point of view, will provide for the child's well-being. Thus a parent who does always what he justifiably believes is the right thing for his children violates no subjective obligation regarding them.

A similar distinction may be made with regard to epistemic obligation. More specifically, we may talk about what are in fact one's obligations with respect to gaining truth and avoiding error, and, alternatively, we may talk about what such obligations are from a particular subject's point of view. The distinction is manifested by the case of two believers, one of whom believes *p* on the basis of what is in fact good evidence, the other of whom believes *p* on the basis of what from his point of view is good evidence. Also in keeping with the moral analog, epistemic responsibility is a function of subjective obligations. Thus an agent is morally blameworthy only if he violates what from his point of view is a moral obligation. As Thomas Nagel points out, an agent who unwittingly does wrong is not responsible for his wrong doing. His action might be a bad thing, and it might even be a bad thing that he exists, but he is not morally responsible for his action unless from his point of view his action is a bad thing.⁵ This is why we do not blame children and psychopaths, although we can make other kinds of moral judgments concerning them. Likewise, epistemic responsibility is a function of subjective obligations. Roughly, a subject believes *p* responsibly if and only if, from his point of view, in believing *p* he is not violating any epistemic obligations (i.e., no obligations with respect to gaining truth and avoiding error).

With this conception of justification in mind, we may now turn to a consideration of some popular forms of internalism.

2. PERSPECTIVAL INTERNALISM

In its broadest formulation internalism is the idea that, in some important sense, the factors which give rise to justified belief must be internal

to the believing subject. Thus internalism in general may be understood as laying down a special sort of requirement for justified belief. Different versions of internalism spell out this requirement in different ways. Externalism, in turn, may be understood as the denial that any internalist requirement on justified belief exists. More specific versions of externalism may be defined as denials that some specific internalist requirement exists.

The first kind of internalism which I would like to consider is what Alston calls “perspectival internalism”. As Alston defines it, this broad theory of justification requires 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 he takes account”.⁶ Thus perspectival internalism in general lays down the requirement that whatever contributes to the justification of a belief must be within the believer’s perspective. Different versions of the theory are defined according to how the subject’s perspective is understood.

The following passages from BonJour indicate that he accepts a version of perspectival internalism, and that he does so on the basis of a responsibilist conception of justification. Arguing against foundationalism BonJour writes,

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

Thus if basic beliefs are to provide a suitable foundation for empirical knowledge, if inference from them is to be the sole basis for the justification of other empirical beliefs, then that feature, whatever it may be, in virtue of which an empirical belief qualifies as basic, must also constitute an adequate reason for thinking that the belief is true. And now if we assume, plausibly enough, that the person for whom a belief is basic must *himself* possess the justification for that belief if *his* acceptance of it is to be epistemically rational or responsible, and thus apparently that he must believe *with justification* both (a) that the belief has the feature in question and (b) that beliefs having that feature are likely to be true, then we get the result that this belief is not basic after all, since its justification depends on that of these other beliefs.⁷

Although the above passage is directly concerned with the possibility

of foundational beliefs, it suggests both a commitment to perspectival internalism and the idea that this commitment is required by taking epistemic justification to be a matter of epistemic responsibility. Bonjour maintains that justified belief just is epistemically responsible belief, and that epistemic responsibility requires that the subject himself “possess” the justification for his belief.

In the above passage Bonjour identifies the subject’s perspective with the totality of her justified beliefs, but we have already noted that the subjective perspective may be more narrowly or more broadly defined. Before raising objections to perspectival internalism in general, it will be helpful to determine its most plausible version, according to the most plausible option for defining the subjective perspective or viewpoint.

As noted above, Alston defines perspectival internalism so as to include three major options, so that the subject’s perspective may be understood as the totality of the subject’s justified beliefs, her beliefs in general, or her knowledge. But the subjective perspective could be defined more broadly still. For it is plausible that in some cases the justification of belief depends on the believer’s perceptual experience, or perhaps some other kind of experience. This is a controversial thesis which has been discussed at length by foundationalists and their opponents, but it is a possibility that we should not ignore. Accordingly, I will understand the subjective perspective defined in the broadest sense to include the subject’s current experience.

We should also ask whether a sensitivity to inference relations should be included in S’s perspective. Thus Alston defends an internalism in which the epistemic agent must have the grounds of her belief safely accessible, but need not be aware of the inference relation between her grounds and the belief in question.⁸ If she accepts her belief on the basis of grounds which in fact support her belief, she need not be aware of this supporting relation in order to come out justified in her belief. It seems incorrect, however, that a subject can be justified on the basis of grounds that she does not see support her belief. Suppose for example that Jane, who is not very good at math, were to base her belief in a complicated theorem on a set of axioms which do in fact support the theorem. But suppose she does so not because she sees the supporting relation, but because she has reasoned invalidly from the axioms to the theorem. My intuition is that she is not justified in her belief that the theorem is correct. As a second example, suppose that Ork is a newly

arrived visitor from Other Earth, a place where the fact that something looks green gives little indication that it is green. We can imagine that on Other Earth the atmosphere often filters light so as to make such reasoning unreliable. Now suppose that Ork believes that the tree in front of him looks green, and concludes on that basis that it is green. In such a case it seems to me that Ork is not justified in his belief. For although he has made a correct inference, coming from Other Earth he could not realize that this is so. In fact, from his point of view the inference is irrational.

We should conclude, therefore, that sensitivity to inference relations can be a contributing factor to the justification of one's beliefs. As such, the perspectival internalist will want to define the subjective perspective so as to include such sensitivity. How this should be done, however, is not immediately obvious. For although we are often "aware" that some set of evidence supports a conclusion, it is hard to say what this awareness consists in. Perhaps an example will clarify the problem.

We said above that Jane is missing something when she believes the mathematical theorem on the basis of certain axioms, but does not see the supporting relation. We may ask ourselves what it is that she is missing. What has to be the case for it to be true that Jane does see the supporting relation? It seems too strong to require that Jane explicitly believe that some inference rule or rules are acceptable. Typically only logicians have beliefs about the deductive rules which we use in our reasoning, and it is agreed on all sides that no one has successfully characterized the rules which govern our non-deductive reasoning. But if we typically do not have beliefs about the rules which govern correct reasoning, how are we to include sensitivity to such rules in the subjective perspective?

I would suggest that although we do not typically have beliefs about such rules, we do countenance such rules in our reasoning. In other words, we follow such rules when we reason, although the way in which we follow them does not involve having beliefs about them. As John Pollock points out, reasoning is something that we do, and as with other things that we do correctly, our doing it involves a procedural knowledge. Procedural knowledge, however, is not propositional, and as such the having of such knowledge does not entail the having of beliefs.⁹

So although we do not typically have beliefs about the rules which govern our reasoning, we do typically countenance certain rules and

not others. But what is it exactly to say that *S* countenances one rule and not another in her reasoning? I have suggested that it is to say that *S* follows the rule, but it is not just that. The rules that we countenance in our reasoning are the rules which we follow when we reason conscientiously. Thus it makes perfect sense to say that someone is reasoning in a way that she does not countenance. We may consider an analogy with hitting a baseball. A good hitter countenances certain rules about hitting, although it is not necessarily the case that he can articulate these rules. Furthermore, he need not follow the rules all the time. It makes sense to say that he knows how to hit, but that he is not doing a very good job of it at the moment. It also makes sense to blame him for not following the rules of good hitting which he countenances. Similarly, most people know how to reason, although they may not do a good job of it all the time, and although they may at times be blameworthy for reasoning as they do.

I suggest, then, that we have a way to broaden the subjective perspective so as to include sensitivity to the rules governing one's reasoning. We may say that *r* is within *S*'s perspective in the broadest sense if and only if *r* is a belief of *S*'s, or a current experience of *S*'s, or a rule of reasoning which *S* countenances. The above discussion has taken place in terms of rules governing reasoning from one's evidence to conclusions supported by the evidence. However, we should recognize the existence of other kinds of rules of reasoning. For example, presumably there are rules relating specifically to the defeat of evidence.¹⁰

We may now turn to an adjudication of the various versions of perspectival internalism. For easier exposition, I will explicitly label the three versions we are considering.

- (PI1) *r* contributes to the justification of *S*'s beliefs only if *r* is a belief of *S*'s, or a current experience of *S*'s, or is a rule of reasoning which *S* countenances.
- (PI2) *r* contributes to the justification of *S*'s beliefs only if *r* is a justified belief of *S*'s.
- (PI3) *r* contributes to the justification of *S*'s beliefs only if *r* is an item of knowledge for *S*.

As Alston notes, restricting justifiers to things which are known by *S* seems too strong. Suppose, for instance, that Jack has all the reason in the world to think that Jill is honest, and is therefore justified in his belief that she is. It seems that on the basis of this belief he is justified

in believing that at least one person he knows is honest. And this is so whether Jill really is honest or not. We may conclude therefore that if any version of perspectival internalism is correct it is either PI1 or PI2.

Alston argues that PII internalism is too weak. Thus if Jack's belief that Jill is honest is wildly irrational, it cannot serve to justify others of Jack's beliefs. He therefore concludes that PI1 may be rejected as an interesting version of internalism. This conclusion, however, is too hasty. For perspectival internalism is not a theory merely about the grounds of a subject's belief, but about any contributor to justification whatsoever. Thus it is more or less uncontroversial to suggest that the grounds of a justified belief must be justified themselves. It is more or less uncontroversial, therefore, that Jack's belief about Jill must be justified if it is to serve as grounds for another belief. But the externalist does not suggest that the causal ancestry of a belief, or the reliability of a belief, contributes to its justification as a ground. The suggestion is that these aspects of a belief may contribute to its justification in a different way. The mere fact that a belief has a certain causal ancestry, for example, is considered to be a factor in the justification of the belief, whether or not S grounds her belief on this fact. If perspectival internalism is to be interesting, therefore, it must be a theory about any contributor to justification whatsoever.

Understood in this way, it is not so obvious that PI1 internalism is too weak. Granted that the grounds of S's justified beliefs must be justified themselves, it is not so obvious that any contributor to justification must be so. For example, the justification of a belief may be affected by counterevidence, or evidence against the belief. Thus if Jack believes that everyone he knows, including Jill, has told lies in the past when it has been to their convenience, this would count against the belief that he knows at least one person who is honest. But what is more, it seems that this belief would affect justification adversely whether or not the belief is itself justified. For if Jack sincerely believes that everyone he knows has lied extensively in the past, he cannot responsibly believe that he knows at least one person who is honest.

In the case of counterevidence, we have an example of something which may contribute to the justification of a belief, but not through being grounds for the belief. As suggested above, sensitivity to the fact that one's grounds support one's belief may also play a role in the justification of belief. And since such sensitivity seldom takes the form of a belief, it seldom takes the form of a justified belief. PI1 therefore

emerges as the strongest version of perspectival internalism so far considered. The central idea of PI2 internalism, that the grounds for one's justified beliefs must be justified themselves, seems correct. However, not all contributors to a belief's justification are grounds for the belief.

Before concluding that PI1 is the strongest version of perspectival internalism, however, we should consider another version suggested by Alston. Alston suggests a "higher-level" requirement for perspectival internalism, requiring not only that whatever justifies a belief be within the subject's perspective, but that the subject be justified in believing that this is so. Alston considers this to be the preferred version of perspectival internalism, although it is a position which he ultimately rejects.

Once again, different versions of this broad position may be defined according to how one understands the subjective perspective. Most of these may be rejected for reasons similar to the considerations above. However, three interesting versions emerge by considering the grounds of S's beliefs rather than contributors to justification in general, and by asking whether there is a higher-level requirement placed on those grounds. Thus we may ask whether one of the following conditions must hold: (1) that S believes her grounds are justified; (2) that S is justified in believing her grounds are justified; (3) that S knows that her grounds are justified.

We may reject the first condition on the ground that it implies that only subjects with the concept of epistemic justification can be justified in their beliefs. This may be seen in that the condition requires that S believe that she is justified in believing her grounds. But in order for S to have such a belief S must have available to her the concepts which go into it. Yet many subjects to whom we wish to attribute justified belief clearly do not fulfil such a requirement. Small children are the most obvious case in point.

Since justified belief in q and knowledge of q each entail belief in q , we can see that conditions (2) and (3) are subject to the same criticism. A defender of either condition might wish to make the following revision in order to save her position. She might stipulate that *if* S believes that she is justified in believing her grounds, then the relevant higher-level restriction applies. Thus with respect to condition (2), the position would be revised as follows: S is justified in believing p on the basis of grounds g only if (i) S is justified in believing g , and (ii) if S believes that she is so justified, then the belief that S is so justified is justified

for S. An analogous revision could be made for condition (3). The problem here, however, is that such revisions are completely ad hoc. Specifically, the internalist cannot explain why a higher-level requirement is necessary for the subject who has the concept of epistemic justification, but not for the subject who does not.

On the basis of the above discussion, PI1 emerges as the preferred version of perspectival internalism. And it seems that it is a strong position indeed. For as BonJour points out, epistemic responsibility does seem to be intimately connected with the believer's subjective perspective on the world. Insofar as we are concerned with a responsibility conception of justification, perspectival internalism *seems* to lay down a legitimate requirement for justified belief.

A not unrelated argument in favor of perspectival internalism is based on a widely held and persistent intuition; namely, that two subjects with the same perspective must have equal justification for any like belief. That is, if S1 and S2 have the same beliefs and experiences, and if they countenance the same ways of reasoning, then they must have equal justification with regard to any like beliefs. But if this is so, then the only thing that can figure into the justification of belief must be beliefs and experiences and the rules which govern reasoning. This intuition is illustrated by the inhabitant of a Cartesian dream world. Whether or not such a subject has knowledge, it seems that he must be justified in believing many of the same things we are. Even if the dream worlder is the victim of a malicious demon when he seems to contemplate a ball of wax in front of a fire, it seems undeniable that he is nevertheless justified in believing that he is sitting before a fire, that he is contemplating a ball of wax, etc. Or at least he is as justified as we are. For from his perspective, there is no difference between the dream world and the real world. For all that he knows, his situation is exactly the same.

In a discussion of internalism, Steven Luper-Foy goes so far as to say that, whatever our analysis of justification, the following "Principle of Irrationality" had better turn out true: "If in one world W1 S's belief that *p* is justified, then it is justified for S in any world W2 that is perceptually and doxastically similar to W1 for S".¹¹ Richard Foley invokes a similar intuition in the following passage.

Suppose that a person is caused to believe *p* by a reliable perceptual process (however this is understood) and that there is available to him no reliable defeating process. According to the reliabilist then, S rationally believes *p*. Now, imagine a second situation

in which S again is caused to believe *p* by a perceptual process, and suppose that in the second situation S thinks, believes, experiences, seems to remember, etc. exactly what he does in the first situation Finally, suppose that in the second situation the perceptual process which causes S to believe *p* is unreliable. Perhaps there is some malfunction in S's optical nerves or perhaps there is an evil demon present

. . . [I]n one sense of rational, [S's] belief in the actual situation is no more and no less rational than it would be in the second situation. What this implies, in turn, is that if there is some factor present in the actual situation which is not present in the hypothetical situation, this factor need not be present in order for the belief to be rational in this sense.¹²

But despite the various considerations in favor of perspectival internalism, I have promised that that position is false. Moreover, I have suggested that the position is not in fact supported by a responsibilist conception of justification. I now turn to arguments to this effect.

3. OBJECTIONS TO PERSPECTIVAL INTERNALISM

In understanding the deontological or responsibilist conception of justification we have made extensive use of the analogy to moral responsibility. By pushing this analogy we may see that perspectival internalism cannot be quite right if the responsibilist conception is correct. For at least two factors which contribute to moral (or epistemic) responsibility may lie outside the subject's present perspective, even when that perspective is interpreted in the broad sense of PI1.

We may consider the first way in which epistemic responsibility may be affected by factors external to the believer's perspective by making use of a Kantian distinction. In the *Groundwork of the Metaphysics of Morals* Kant says the following.

For if any action is to be morally good, it is not enough that it should conform to the moral law – it must also be done *for the sake of the moral law*: where this is not so, the conformity is only too contingent and precarious, since the non-moral ground at work will now and then produce actions which accord with the law, but very often actions which transgress it.¹³

Perhaps we may make the same (or at least a similar) distinction by saying that while some actions are merely *in accordance* with our (subjective) moral obligations, others of our actions are *in conformance* with our (subjective) obligations. The former do not conflict with our obligations because, by happy chance, the non-moral motives for our action do not create such a conflict. The latter actions are of a superior

status. They do not fall short of our obligations because their motivation includes consideration of those very obligations. When our actions are of the latter kind, we act (at least partly) because this is what our obligations require. It is clear that only the latter actions merit praise. While it is generally a good thing that our actions are in accordance with our obligations, we are morally praiseworthy only when they are in conformance with our obligations. Thus moral responsibility involves conformity to one's obligations.

But now the same distinction can be made with respect to our beliefs. While some of our beliefs are merely in accordance with what we take to be our epistemic obligations, others of our beliefs are in conformance with the epistemic norms which we countenance. The latter beliefs are accepted (at least partly) because we countenance certain epistemic norms.

To illustrate the distinction we may consider the following case. Suppose that Larry wants to balance his checkbook, and correctly believes on the basis of a lucky guess that his outstanding checks amount to a total of \$200.00. Suppose also that Larry knows the amount of each outstanding check but has not bothered to calculate the total. He could perform the calculation just on reflection, but fails to do so. Now in this situation the epistemic norms which Larry countenances would permit the belief that the total of outstanding checks is \$200.00, since those norms include the simple rules of addition involved in the required inference. But although Larry's belief is in accordance with the norms, it is not in conformance with the norms, for his belief does not utilize the norms in any way. He does not believe that the total is \$200.00 because he countenances an inference to that effect. As a result, Larry is not justified in believing that his outstanding checks total \$200.00.

But notice that whether Larry's belief is in conformance with the epistemic norms he countenances, or whether it is merely in accordance with those norms, is not typically a fact which is within Larry's perspective. He will normally have no beliefs whatsoever concerning this state of affairs. Therefore perspectival internalism is false, for there is at least one factor which contributes to the justification of belief, but which need not be within the believer's perspective.¹⁴

Another way in which factors outside one's perspective may contribute to the justification of belief is through the violation of epistemic norms at an earlier time. Consider, for example, a person who reasons his way to a conclusion through an extended series of inferences. We

can imagine that the series is long enough so that inferences made at the beginning cannot be kept in mind in the latter stages of the process. We can imagine, in other words, that earlier inferences on which later conclusions depend are no longer within the subject's perspective. But of course mistakes in those earlier inferences will nevertheless affect the justification of the later conclusions. Here again, the justification of a belief depends on some factor outside the perspective of the believer.¹⁵

Or suppose that Maria believes that her favorite singer is Italian. She believes this because she seems to remember clearly that it is so, and she presently has no reason for doubting her belief. Thus let us assume that Maria's belief is in perfect conformance with the epistemic norms which she countenances, and to that extent her belief is perfectly responsible. Assume also, however, that Maria first came to her belief on the basis of testimony from her mother, who believes that everyone she likes is Italian. At the time Maria knew that her mother was an unreliable source in these matters, and she realized that it was not rational to accept her mother's testimony. However, she believed anyway.

Clearly Maria is not justified in her belief now, even though her belief is presently in conformance with the epistemic norms which she countenances. The reason is that her belief involves epistemic negligence at an earlier time. Maria ought not to have believed that her favorite singer is Italian in the first place, so she cannot become justified by seeming to remember that he is now. If this were not the case, then justification would come quite easily to those who are completely irresponsible in their belief formation, but who also have a propensity to forget how they first arrived at their beliefs. But of course we do not say that such beliefs are justified. This conclusion is in keeping with the analogy to moral responsibility. It is widely accepted that one cannot escape moral blame for an action merely by doing what one thinks is right now. For it might be the case that one ought to know better.¹⁶

The upshot of the last two examples is that we can presently fail in our responsibilities through earlier negligence. And of course whether a belief involves earlier negligence is not typically within the believer's cognitive perspective. Such transgressions can go unnoticed or they can be forgotten. And so we have another way in which factors outside one's perspective may contribute to the justification of one's beliefs.

It should be noticed that the above discussion also undermines the other main consideration we have mentioned in favor of perspectival internalism – namely, the intuition that two subjects with the same perspective must be equally justified in like beliefs. It does so by showing just where that intuition goes wrong. It is now apparent that two subjects with the same perspective need not have arrived at their beliefs in the same way. Thus it might be that one subject has reasoned carefully to a certain belief while the other has formed the same belief on a whim. Or it might be that two subjects have the same present beliefs and apparent memories, but that one has his memories infected by past violations of his epistemic obligations while the other enjoys a history of unwavering rationality. Thus neither of the considerations most widely offered in favor of perspectival internalism will support that theory.

I now turn to a discussion of two other popular kinds of internalism.

4. ACCESS INTERNALISM AND LEVEL INTERNALISM

We defined perspectival internalism as requiring that all contributors to the justification of a belief be within the believer's cognitive grasp or perspective. An alternative manner of restricting the class of justifiers is that preferred by Alston, and amounts to restricting justifiers to items to which the believer has specially favored access. Thus what is important is not whether a justifier r is within S's cognitive grasp, but whether S has some privileged access to r , such that r could come properly within S's grasp. As Alston suggests, access internalism is weaker than perspectival internalism. Whereas perspectival internalism restricts justifiers to what is already within the subject's cognitive perspective, access internalism restricts justifiers to what can be brought into that perspective, via some special form of access.

In the following discussion I will consider a version of access internalism offered by Ginet, along with his argument that the position is supported by the responsibilist conception of justification. In Chapter III of *Knowledge, Perception and Memory*, Ginet lays down the following definition of justified belief: "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 ".¹⁷ Thus Ginet's conception of epistemic justification is a deontological one; the central idea is that the believer stay within his epistemic

obligations, that he not be reproachable for his belief. A few pages later Ginet gives us his version of access internalism. He writes,

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.¹⁸

What is especially interesting for our purposes is that Ginet thinks that the requirement laid down in the above passage is supported by his (deontological) conception of justification. His argument seems to be this: Epistemic justification is a matter of epistemic responsibility; it is a matter of what the subject ought and ought not to believe. Furthermore, such obligations imply that the subject be capable of fulfilling them; this is simply a matter of the standard assumption that "ought" implies "can". But the epistemic subject cannot fulfil these obligations unless the conditions for doing so are directly recognizable to the subject, or are entailed by what is directly recognizable to the subject. Therefore epistemic justification, conceived as a deontological concept, implies the restriction on justifying conditions laid down by access internalism.¹⁹

Ginet seems to be offering support for his last premise in the passage that immediately follows.

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 clear-headedly and attentively he considers his position he will detect no change. If it seemed before to S that he had justification for being confident that *p* then it must still seem so to him.²⁰

This last passage reads very much like the one quoted from Foley above. Like Foley, Ginet seems to be insisting that two subjects with the same inner perspective must be equally justified in holding like beliefs. So we see that arguments similar to those offered in favor of perspectival internalism have also been invoked in favor of access internalism.

But of course access internalism is false for the same kinds of reasons that perspectival internalism is false. The requirement laid down by

access internalism is too strong in that part of what makes a belief justified may not be directly accessible to the believer. First, whether a belief is presently in conformance with proper epistemic norms, as opposed to being merely in accordance with such norms, may be a fact which is not directly accessible to the believer. The difference between these two states of affairs is partly a function of the causal processes which lead to S's belief, but typically the relevant facts concerning such processes will not be directly accessible to S. Secondly, instances of prior negligence are typically not directly accessible to a believer. But as we saw above, past mistakes can effect present justification.

Thus it is false that everything which contributes to the justification of belief must be something which is directly accessible to the believer. The requirement laid down by access internalism is too strong. We may see by the following set of examples that the requirement is also too weak.

Suppose that Harry believes on a whim that someone has recently attempted to assassinate the President. Therefore he is not justified in believing as he does. But suppose also that lying unnoticed on the table in front of Harry is a newspaper proclaiming the event in bold headlines. If Harry were to simply notice the newspaper, he would be justified in believing that someone has attempted to assassinate the President. Here Harry has easy access to sufficient justification for his belief, but his belief is not justified. The reason that we conclude that this is so is because Harry has not used the accessible, but merely potential, justification. This being so, his belief remains unjustified and irresponsible.

Of course the access internalist has an easy response to the example. For it will be argued that an appropriate version of access internalism restricts justifiers to that which may be "internally" accessed. Thus Alston restricts justifiers to that which can be fairly readily accessed, "just on reflection", and Ginet restricts justifiers to what is "directly recognizable". But the example serves to point out what is wrong with access internalism: potential justification is of no importance if it is not used by the believer. The fact that a justifier is readily accessible seems irrelevant. Other examples can be constructed so that the potential justification is internally accessible, and can be used to drive home the same point.

Thus suppose that Wendy is a witness to a robbery. In court she testifies that there were four robbers. However she believes this not on

the basis of memory, but on the basis of an unjustified belief of hers that robbers always work in groups of four. Suppose also that if Wendy were to review in her mind the event of the robbery, she would easily remember that there were in fact four robbers. Again, we must conclude that the believer does not have justified belief, even though she does have readily accessible justification for her belief. Unless her potential justification is used, her belief is epistemically irresponsible and therefore unjustified.

The above example suffices to show that access internalism is too weak, at least in so far as it is a theory about deontological justification. An obvious requirement of epistemic responsibility is that a subject use her evidence, not that she merely have evidence available.

During his discussion of access internalism, Gettier hints that he accepts yet a third kind of internalism. He says, "assuming that he has the relevant concepts, S can always tell whether or not he has justification for being confident that *p*".²¹ One might interpret this sentence as the claim that, assuming he has the relevant concepts, if S is justified in believing something, then S can know that he is justified in believing that something. I understand by 'level internalism' any position which claims that in order for a belief to enjoy some important epistemic status for S, the belief that it does enjoy that status must enjoy an important epistemic status for S. Thus, as the name implies, this kind of theory holds that justification or knowledge for one level of belief implies justification or knowledge for a higher level of belief. Gettier explicitly endorses a version of level internalism later in the same chapter when he endorses the following claim: "If I know that *p*, then I can, merely by attentive and clear-headed reflection (provided that I understand the proposition that I know that *p*), directly recognize that I am justified in claiming to know that *p* . . .".²²

In *The Foundations of Knowing*, Chisholm cites H. A. Prichard as a defender of level internalism. Thus Prichard writes, "Whenever we know something, we do, or at least can, by reflecting, directly know that we are knowing it".²³ Perhaps the most explicit defender of level internalism, however, is Chisholm himself, who has endorsed several versions of the theory. Thus Chisholm writes, "Speaking somewhat loosely, we may say that, if a proposition is reasonable for a given subject, then it is reasonable for that subject to believe the proposition to be reasonable . . . Hence reasonability is properly called an 'internal' concept".²⁴ Chisholm has also accepted stronger versions of level in-

ternalism, for example when he has taken the following proposition to be an axiom: If *S* is justified in believing that *p* and *S* believes that she knows that *p*, then *S* is justified in believing that she knows that *p*.²⁵

I am not aware of any explicit argument that level internalism is supported by a responsibilist conception of justification. But I have suggested that there is the hint of such an argument in Ginet, and I have noted that the level internalist Chisholm has strongly endorsed a responsibilist conception. Therefore even if no philosopher has explicitly suggested a supportive relation between the two positions, it will be interesting to investigate whether such a relation exists.

We may sufficiently argue against all varieties of level internalism by arguing against its weakest version; *S* is justified in believing *p* only if *S* is justified in believing that *S* is justified in believing *p*. For if we assume that knowledge of *p* implies justified belief in *p*, then what we say about the weakest version will apply equally well (or with slight revision) to stronger versions.²⁶ Let us symbolize the present version of level internalism as follows.

(J) $Jp > JJp$

We may now turn to arguments against (J).

The first objection which I would like to raise is that many epistemic subjects simply do not have beliefs about the justification of their beliefs. However, we should not conclude that therefore such subjects do not have justified beliefs. For example, *S* might not have any belief about whether she is justified in believing she exists. But we should not say that therefore she is not justified in believing that she exists. But if (J) is true, we are forced to this conclusion. For according to (J), $-JJp$ implies $-Jp$.

Another objection turns on the fact that, for any proposition *p* which *S* is justified in believing, (J) requires not only that *S* is justified in believing that *S* is justified in believing *p*, but also that *S* is justified in believing that *S* is justified in believing that *S* is justified in believing *p*. For by a second application of (J), we see that JJp implies $JJJp$. Further, a third application of (J) implies $JJJJp$, and so on. It is unreasonable, however, to assume that epistemic subjects actually have such an infinite hierarchy of beliefs. It is implausible, for example, that I can even grasp the proposition indicated by ‘ $JJJJp$ ’.²⁷

The above considerations show that (J) is false. However, the level internalist may revise his position so as to avoid the present objections.

Thus he may revise his position as follows: if S is justified in believing p and S believes that she is justified in believing p , then S is justified in believing she is justified in believing p . A look at the passages from Chisholm and Ginet above will show that they in fact use similar formulations to represent their positions. An important objection to such a revision, however, is that it seems ad hoc. Why should second-level justification be required of subjects who believe that their beliefs are justified, but not of those who do not? If second-level justification is required in some cases it is hard to understand why it is not required in all of them.

Even if such an explanation can be provided, there are further objections which may be leveled against the new formulation. For we may imagine that some subjects do not have the reasoning capacity necessary to reach higher levels of justification. Thus suppose Mary is justified in believing there is a tree in front of her and suppose also that she believes she is justified in believing there is a tree in front of her. Now presumably it takes more intricate reasoning to arrive at justification for the belief that one is justified in believing there is a tree in front of one than it does to arrive at justification for the belief that there is a tree in front of one. In other words, it takes more intricate reasoning to arrive at Jp than it does to arrive at p . But then it is not difficult to imagine a situation in which Mary is perfectly capable of arriving at justification for her belief about the tree, but in which it is completely beyond her capacities to arrive at justification for her belief that her belief about the tree is justified. But in that case the new version of level internalism is false. For we have described an example where Mary is justified in believing p , she believes that she is justified in believing p , but she is not justified in believing she is justified in believing p .

At this point the level internalist may wish to make a further revision. Thus he might contend that his position is best represented as follows: S is justified in believing p only if, if S believes she is justified in believing p and S has the reasoning capacity to arrive at justification for her belief that she is justified in believing p , then S is justified in believing she is justified in believing p .

As we shall see below, there are convincing reasons to reject this latest formulation. Another question which arises, however, is whether the latest formulation is the best way to account for the intuitions behind level internalism. I will argue that it is not.

One way to approach the problem is to ask, “Why would someone be a level internalist in the first place?” Chisholm has defended versions of level internalism with the following claim; whatever justifies a subject in believing p is available to the subject to justify her in believing that she is justified in believing p . Thus after considering how it is that a subject can be justified in believing that she is justified in believing a certain proposition, Chisholm says “I suggest that the substrate for this higher order epistemic state is the same as that for the lower order epistemic state which is its objective. In other words, those nonnormative states (the perceptual takings) that justify me in believing that sheep exist also justify me in believing that I am justified in believing that sheep exist”.²⁸ Chisholm’s remarks imply that his motivation for being a level internalist is his adherence to perspectival internalism, or perhaps access internalism. More specifically, he believes that level internalism is true because he assumes that perspectival (or access) internalism is true. He assumes that *whatever justifies S’s lower level belief is available to her*, and therefore available to justify her higher level belief.

This contention is confirmed by the following passage from Chisholm.

Now, I think, we may characterize the concept of “internal justification” more precisely. If a person S is *internally justified* in believing a certain thing, then this may be something he can know just by reflecting upon his state of mind. And if S is thus internally justified in believing a certain thing, can he also know, just by reflecting upon his state of mind, that he *is* justified in believing that thing? This, too, is possible – once he has acquired the concept of epistemic justification.²⁹

According to Chisholm, we can know whether we are justified in believing a proposition simply by reflecting on our state of mind. But this could be so only if everything which contributes to the justification of a belief is within the believer’s perspective, or available to that perspective just on reflection. In other words, it could be so only if perspectival or access internalism is true. And this, of course, is the reason he also accepts level internalism. Because he thinks that everything which contributes to the justification of a belief is available to the believer, and because he thinks that the same things which justify p for S are also sufficient to justify Jp for S (provided that they are available), Chisholm believes that first-level justification implies second-level justification.

The motivation for level internalism, then, seems to be the acceptance of perspectival or access internalism. But is level internalism true?

If it is a theory about the responsibilist conception of justification then it is clear that it is not true. For like access internalism, level internalism fails to require that the subject use her evidence or reasoning capacities in order to arrive at her belief. We may imagine a case where, although S believes she is justified in believing p , and although she has the evidence and reasoning capacity necessary to arrive at justification for her second-level belief, S does not use her evidence to arrive at her belief. For example, we may imagine that S believes she is justified in believing p not because she has reasoned correctly to this conclusion, but because she wants to be accepted into the social circle of some anti-skeptical epistemologists. It is clear that in such a case S's belief is epistemically irresponsible and therefore unjustified in the deontological sense.

5. NORM INTERNALISM

The last two sections suggest the following characterization of justified belief, which is in accordance with a responsibilist conception of justification.

- (NI) S is justified in believing p iff S's believing p is in conformance with a correct set of epistemic norms (rules of reasoning), and the history of S's belief has also been in conformance with those norms.

In the remainder of the paper I will try to fill out this position, first by way of some general remarks, and then by comparing the position with others which are close to it. Of course a major task will be to explicate what makes a set of epistemic norms "correct".

First, it will be noticed that in cases of indirect (mediate) justification, (NI) requires that S use the grounds which she has for her belief, if that belief is to come out justified. This is because conformance to a set of epistemic norms requires that one's belief be guided by the norms in the set. Thus for cases of indirect justification, (NI) ensures that S believe p because p is supported by evidence which S has.

Secondly, by the history of S's belief p I mean to include the genesis of p , as well as the genesis of the beliefs which p depends on for its justification. Thus our account requires not only that S's present belief be responsible, but that the evidence for S's belief be responsible as well. I take it that this requirement is in accordance with our intuition

that evidence must itself be justified in order to give rise to further justification.

We should also try to become more clear with respect to just what epistemic norms are. We have described such norms as rules of reasoning, or rules of belief acquisition and maintenance. Further, it is plausible that the antecedents of these rules must be in terms of states of affairs which are in an important sense directly accessible to the cognitive faculties of the believer, so that the rules can be used in the guidance of belief acquisition and maintenance. Thus we may follow Goldman in conceiving epistemic norms as functions “. . . whose inputs are certain conditions of a cognizer – e.g., his beliefs, perceptual field, and ostensible memories – and whose outputs are prescriptions to adopt (or retain) this or that doxastic attitude – e.g., believing *p*, suspending judgment with respect to *p*, or having a particular subjective probability vis-à-vis *p*”.³⁰

We can now ask with Goldman, What makes a certain [set of epistemic norms] right or correct? Another way to ask the same question is to ask, What makes a set of epistemic norms the norms which give rise to justification? Perhaps the best way to investigate this question is to take a look at what some others have said.

Goldman’s answer to his own question is best understood in two parts. Remember that his term for a complete set of epistemic norms is “doxastic decision principle”, or “DDP” for short. Goldman writes,

Assume that a unique DDP is correct or right. Then *S* is justified in believing *p* at *t* if and only if the right DDP, when applied to the relevant conditions that characterize *S* at *t*, yields as output the prescription “believe *p*”.³¹

He then goes on to say,

We can indicate the sort of condition being sought by formulating the condition appropriate to *externalism*, viz, (1):

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

By “optimal” I mean, of course, optimal in producing true belief and error avoidance.³²

Goldman’s answer to the above question is straightforward; the right set of epistemic norms is that which is optimal in gaining truth and avoiding error. These norms give rise to justification when a believing subject adopts the right belief under the right conditions. But notice that in the above quotation Goldman does not require that the believer use the norms in arriving at her belief. All that is required is that the

believer fall under application of the rules involved. If the believer satisfies the appropriate antecedent and assumes the appropriate belief, then the believer is justified in her belief. No requirement is mentioned to the effect that the believer must be guided by the norms in arriving at her belief.

I would argue that Goldman's view can be criticized on two counts. First, epistemic justification requires that belief be in conformance with the relevant epistemic norms, not merely in accordance with such norms. Secondly, the right epistemic norms should not be characterized in terms of actual optimality. For the norms which are actually optimal might not be the norms which a believer countenances. Imagine that we are inhabitants of a Cartesian dreamworld. In such a world our sensory beliefs would be highly unreliable, and therefore the norms which we countenance in the acquisition and maintenance of such beliefs would also be unreliable. But surely we would nevertheless be justified in many of our sensory beliefs. The view I will wish to defend is that epistemic justification depends essentially on the norms which we in fact follow when we are reasoning conscientiously, not on the norms which are actually optimal (although these may in fact be the same norms).

Goldman has argued that whether a set of epistemic norms is correct depends on whether it is optimal, so that correctness is a contingent matter. I have also argued that whether a set of epistemic norms is correct is at least in part a contingent matter, for whether a set of norms gives rise to epistemic justification depends on whether it is the set which is in fact countenanced by a given believer. It might be argued, however, that if a set of epistemic norms is correct then this is necessarily so. Thus Pollock has argued that correct norms are analytically correct. And although Chisholm does not speak of epistemic norms specifically, a natural extension of his views would be that the correctness of such norms is synthetic a priori.

Pollock recommends his view at least partly on the ground that it avoids epistemic relativism. It is not obvious to me, however, that conceiving of epistemic norms as necessarily correct does solve the problem of relativism. For there might be more than one set of epistemic norms which is necessarily correct. This last possibility is especially evident with respect to Pollock's view. I will quote him at length.

Let us take the *conceptual role* of a concept to consist of (1) the reasons (conclusive or

prima facie) for thinking that something exemplifies it or exemplifies its negation, and (2) the conclusions we can justifiably draw (conclusively or prima facie) from the fact that something exemplifies the concept or exemplifies the negation of the concept. My proposal is that concepts are individuated by their conceptual roles. The essence of a concept is to have the conceptual role that it does

The importance of this theory of concepts for the matters at hand is that it lays to rest the spectre of epistemological relativism. Epistemological relativism is the view that (1) different people could have different epistemic norms that conflict in the sense that they lead to different assessments of the justifiedness of the same belief being held on the same basis, and (2) there is no way to choose between these norms. The epistemological theory of concepts enables us to escape any such relativism. Because concepts are individuated by their conceptual roles, it becomes impossible for people's epistemic norms to differ in a way that makes them conflict with one another. The epistemic norms a person employs in reasoning determine what concepts he is employing because they describe the conceptual roles of his concepts. If two people reason in accordance with different sets of epistemic norms, all that follows is that they are employing different concepts.³³

Pollock's view does avoid epistemic relativism as he defines it, but it clearly does not escape other kinds of relativism which might seem as problematic. Thus the view does not preclude that justification may arise through different sets of epistemic norms. And although it provides a basis for choosing which norms are applicable to a given belief, it does not provide a basis for choosing what concepts and therefore what beliefs to employ.

The view seems to be objectionable for other reasons as well. For it at least seems to imply that two people cannot genuinely disagree about whether a body of evidence supports a conclusion. For any ostensible disagreement would indicate that they are employing different concepts and therefore talking about different conclusions and different evidence. But genuine disagreement of this sort seems commonplace.

Alternatively, one might hold that the correctness of epistemic norms is synthetic a priori. Thus Chisholm holds that the following epistemic principle is a synthetic a priori truth.

- (B) For any subject S, if S believes, without grounds 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.

Although (B) is not an epistemic norm, and Chisholm does not speak about epistemic norms directly, it is only natural that someone who thinks (B) is a synthetic a priori truth would also think that it is a synthetic a priori truth that a corresponding epistemic norm is correct.

But of course the central issue here is not whether Chisholm endorses or would endorse such a position, but whether the position is correct.

One reason that such a position seems false applies to any position which holds that correct epistemic norms are necessarily correct. The idea is that it is possible for different cognitive subjects to be constructed very differently, so that the ways they go about getting information may be very different as well. As such, these different ways will be governed by different norms. But if correct epistemic norms are necessarily correct, then it will be possible for subject A to get justified beliefs via the norms of subject B, even though subject A does not countenance B's norms. An example should help to illustrate the problem.

Suppose that a certain species of Martian is constructed with very poor visual apparatus, but is highly skilled at acquiring information in some other way. As a result, these Martians are born with no strong propensity to formulate beliefs about the external world on the basis of what they seem to see. In other words, nature does not wire them with the norms governing visual beliefs that we now employ so naturally. But suppose that a particular Martian irresponsibly formulates a belief on the basis of what he seems to see, and that this belief is in accordance with the norms that now govern our visual belief acquisition, perhaps with a norm corresponding to some vision-specific alternative to (B). By hypothesis the belief is irresponsible – it is not in conformance with the norms which the Martian countenances. Is the belief nevertheless justified? Someone who thinks that our perceptual norms are necessarily correct would have to say yes. But in so far as the belief is epistemically irresponsible it cannot be justified.

Thus none of the views we have considered seems correct. One problem which all of the views have in common is that they do not require that the norms which give rise to justification are countenanced by the believing subject. On the other hand it seems false that just any rules of reasoning can give rise to justified belief, if only they are adopted by the relevant believer. This suggests the following position: We should retain the idea that epistemic norms must be countenanced in order to give rise to justification, but we reject the idea that any norms will do. We may spell out this position more specifically as follows.

First, in order for a set of norms to give rise to justification for S, S must countenance those norms. In other words, the norms must be

those which are used by S when S reasons conscientiously. This view arises naturally from the conception of epistemic justification as epistemic responsibility. By saying that the norms which give rise to justification must be those which are countenanced by S, we point to the idea that S must reason in a way that is right from S's point of view.

But can S reason in any way? We have said that for epistemic norms to give rise to justified belief, they must be countenanced by the believing subject. But we should not say that this is sufficient for correctness of epistemic norms. We should not say that any set of epistemic norms, if countenanced, would be correct. Thus we can imagine that some people, even when they reason conscientiously, reason superstitiously. It is a bit harder to imagine, but perhaps possible, that some people conscientiously follow the norm 'Believe what makes me happy', or 'Believe what comes to mind'. We will want to rule out the possibility that these kinds of norms can be used to arrive at justified belief.

One way in which we can rule out such norms is to stipulate that any complete set of epistemic norms must be somewhat extensive. In this way an occasional bad rule will often be overridden by good rules. For example, even superstitious people typically countenance epistemic norms which would preclude them from being justified in their superstitious beliefs. Thus consider the person who strongly believes in astrology, so that she accepts certain beliefs about the future on the basis of her investigation of the positions of the stars. Although these beliefs might be allowed by some of the rules she truly countenances, they will be disallowed by others. More specifically, they will be disallowed by those governing criteria for acceptable explanations, and those governing defeaters of explanations. Similarly, it is plausible that some people reason according to the gambler's fallacy even when they are being conscientious. Yet these same people will typically countenance other norms as well, by which their beliefs about gambling could be criticized.

But although people will typically countenance good rules which override bad rules, it is not obvious that all people do, or that any believer must. But we can build such a requirement into (NI) if we accept a strong understanding of 'epistemic norm' or 'rule of reasoning', and then agree with Pollock and Chisholm that some epistemic norms are necessarily correct if countenanced by the believer, while others are necessarily incorrect.

(NI) contains the claim that S is justified in believing *p* only if S's

believing p is in conformance with a relevant set of epistemic norms. But a strong understanding of epistemic norms allows us to rule some norms into the set and some norms out on a priori considerations. For it would seem that the very idea of an epistemic norm requires that such a norm have as its purpose the pursuit of truth and the avoidance of error. But it is hard to see how a norm such as 'Believe what makes me happy' could be consistent with such a purpose. Similarly, although an epistemic norm may not in fact be reliable, in order for it to qualify as an epistemic norm there must be a pretense of reliability. But again, norms like 'Believe what makes me happy' and 'Believe what comes to mind' do not carry with them even a pretense of reliability. I suppose that we can imagine an intricate cluster of assumptions for some S , such that those assumptions impose a pretense of reliability on the above norms, but it is hard to imagine that such assumptions would not conflict with other norms which will necessarily be in a correct set. This last claim is supported by the idea that epistemic norms are rules of reasoning. But the very idea of reasoning excludes some doxastic practices from qualifying as such, while it requires that other practices be employed.

We can imagine, I think, beings which reason in ways which are very different from ours. But we would not count just any practice of belief acquisition and maintenance as reasoning. Thus we would not count a practice which gave rise to many and obvious contradictions. Neither would we count practices which were consistently and obviously unreliable. But it is plausible that the conceptual restraints imposed by the idea of "rules of reasoning" go even further than this. I would like to suggest that the very idea of reasoning incorporates the idea of reasoning in a way consistent with *modus ponens*, and in a way consistent with other central aspects of our way of reasoning. In order for a practice of belief acquisition and maintenance to even count as reasoning, it has to be at least consistent with the way we reason, in the sense that it not lead to, on a grand scale, what we would consider wild and irrational results. If this is right then the norms considered above will not count as correct epistemic norms, since they will not even count as epistemic norms.

Let us agree that all epistemic norms, if they are to give rise to justified belief, must be countenanced by the believer. Let us also agree that, once countenanced, some norms will be a priori correct. Are all correct epistemic norms, if part of a countenanced set, a priori

correct? Before answering this question it will be useful to distinguish between initial and adopted norms. Initial epistemic norms are norms which a subject starts with. By using these norms, a subject may eventually adopt others. Thus I now countenance certain norms governing belief about temperature whose antecedents describe perceptual experiences of thermometers. The norms say that given certain experiences of thermometers, and if there are no grounds for doubt, I may believe that the temperature outside is such and such. It is obvious that these norms were not wired into me before birth. I adopted these norms only after learning about thermometers. But it is also obvious that any such norm, if correct, is not necessarily correct. It might have been the case that I did not learn properly about thermometers, that I countenanced the norm hastily or otherwise irrationally. We should conclude, therefore, that adopted norms are restricted by the norms which a subject already countenances. They may be admitted into a correct set of epistemic norms only if their admission is permitted by the norms already in the set.³⁴

Are all initial correct epistemic norms a priori correct? I have already mentioned the reason that I think we should say no. Above we considered Chisholm's epistemic principle (B), and we commented that it does not seem to carry the weight of necessity. Or consider other plausible candidates for perceptual norms. Does it follow from the concept or nature of reasoning that, given a certain tree-like experience *E*, *S* is justified in believing that there is a tree before her? Consider that the experience which a bat has of a tree is very different. But we could have been built with the perceptual faculties of a bat, in which case we would probably have been wired with the following norm; Given perceptual experience *E'* (where *E'* is the experience which a bat has of a tree), it is permissible to believe that there is a tree in front of one. But surely this norm does not follow necessarily from either the concept or the nature of reasoning. After all, we are reasonably familiar with both the concept and the nature of reasoning, but we have no idea what the content of the antecedent of this norm is like. It seems to me that we countenance many of the norms that we do simply because that is the way we were built. If, for example, our perceptual faculties had been built differently, nature would have wired us with different norms. And yet this would not have precluded us from reasoning, or from having justified perceptual beliefs.

Furthermore, we need to say that such norms are initial rather than

adopted. It seems implausible that we adopt such norms only after some kind of induction or other indirect justification process. First, the average believer never considers anything like such an indirect justification, nor is it plausible that one is commonly worked out below the conscious level. Second, such powers of indirect reasoning that would be required are presumably mastered only at a well developed stage, if at all. Yet we want to attribute perceptual knowledge to small children and other unsophisticated believers.

The problem is that many of the norms we need to account for justification cannot plausibly be said to follow necessarily from the idea or nature of an epistemic norm. It seems plausible that while some rules of reasoning follow necessarily from the idea or nature of reasoning, and while other practices are ruled out by these, some rules of reasoning are neither necessarily correct nor necessarily incorrect. We are left with the idea that correct sets of epistemic norms may differ, though they cannot differ in just any way, and they in fact will necessarily be much alike.

Have we been led to an unpalatable relativism? I do not think that we have. Rather, we seem to have arrived at the common sense way of viewing things. That is, people reason in ways which are somewhat different, and other species of epistemic agents might reason in ways which are quite different. All these different practices might lead their practitioners to justified belief. Yet you cannot get justified belief just any way. The moral analog of this situation would be a number of moral agents whose moral practices are somewhat different, although each is inwardly coherent and extensive enough to constitute an alternate moral position. Surely these different moral agents could be morally responsible in their conscientious actions. Yet not any practice is consistent with morally responsible actions.

The present account of justification may be said to be internalist in two ways. The first is that the antecedents of epistemic norms which give rise to justification must be formulated in terms of conditions which are directly accessible to S in the following sense; the cognitive faculties of S must be able to take account of the conditions without S first having to make a judgment about them. Otherwise the norms could not be efficacious in the guidance of S's belief acquisition and maintenance.³⁵ As a result, the norms will be formulated wholly in terms of S's doxastic states (beliefs) and non-doxastic states of awareness, along with the properties of such states that are also directly accessible to S,

if any. Thus the fact that the norms must be used by S leads to the condition that they must be usable. And this leads to internalist restraints on the antecedents of the norms. The second way in which the account is internalist is that the norms involved in justification must be countenanced by S. This condition arises from the fact that, in an important sense, the cognitive practices which give rise to justification must be "right from S's point of view". It is not enough that S's beliefs are in conformance with some set of epistemic norms otherwise external to S. Perhaps this weaker kind of conformance could make S's beliefs highly reliable, but it could not render them responsible.

6. JUSTIFICATION AND KNOWLEDGE

One might agree that (NI) is a correct explication of the responsibilist conception of justification, and yet still wonder whether this conception of justification is that which is involved in knowledge. The problem is that knowledge requires that belief be the result of a reliable cognitive process, but responsibility does not insure reliability.

Perhaps the most central problem with reliabilism is that it allows that a subject may be justified in her belief even when that belief is epistemically irresponsible. The major concern which has been raised with respect to (NI), on the other hand, is that on that theory a subject can have justification, and thus presumably knowledge, without reliability. Both problems can be solved if we make reliability a condition of knowledge, but not of justification.

Such an analysis would recognize that knowledge is virtuous in two ways. First, belief which is knowledge is justified in the sense that it is epistemically responsible. Secondly, belief which is knowledge arises from a reliable cognitive process. Reliability, then, is an additional condition on knowledge. Finally, the two kinds of virtue are related in the following way; in cases of knowledge, it is through conformance to correct epistemic norms that S manifests reliability with respect to believing *p*.³⁶

NOTES

¹ This label was coined by William Alston in his 1985, 'Concepts of Epistemic Justification', *Monist* 68, apparently on advice from Alvin Plantinga.

- ² Roderick Chisholm: 1977, *Theory of Knowledge*, 2nd edn., Prentice-Hall, Englewood Cliffs, New Jersey.
- ³ Laurence Bonjour: 1985, *The Structure of Empirical Knowledge*, Harvard University Press, Cambridge, Massachusetts, p. 8.
- ⁴ The above conditions for moral and epistemic responsibility assume that S is a moral agent and that S is an epistemic agent, respectively. In other words, they assume that S has obligations of the relevant kind.
- ⁵ Thomas Nagel: 1982, 'Moral Luck', in Gary Watson (ed.), *Free Will*, Oxford University Press, Oxford. See also Michael J. Zimmerman: 1987, 'Luck and Moral Responsibility', *Ethics* and his 1988, *An Essay on Moral Responsibility*, Rowman and Littlefield, Totowa, New Jersey.
- ⁶ William Alston: 1986, 'Internalism and Externalism in Epistemology', *Philosophical Topics* XIV, 179–80.
- ⁷ Laurence Bonjour: 1980, 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy* V, 55.
- ⁸ William P. Alston: 1988, 'An Internalist Externalism', *Synthese* 74, 265–83.
- ⁹ For an interesting discussion of this issue see John Pollock: 1986, *Contemporary Theories of Knowledge*, Rowman and Littlefield, Totowa, New Jersey, pp. 124–32.
- ¹⁰ Perhaps a better name for a rule of reasoning is Pollock's 'epistemic norm', since the latter term does not suggest that all belief acquisition and maintenance is the result of inference from other beliefs. Below I will follow both usages.
- ¹¹ Steven Luper-Foy: 1988, 'The Knower, Inside and Out', *Synthese* 74, 361.
- ¹² Richard Foley: 1984, 'Epistemic Luck and the Purely Epistemic', *American Philosophical Quarterly* 21, 113–14.
- ¹³ Immanuel Kant: 1964, *Groundwork of the Metaphysics of Morals*, trans. H. J. Paton, Harper and Row, New York.
- ¹⁴ Here I am indebted to Ernest Sosa.
- ¹⁵ It should be noted that not all perspectival internalists make justification a function of the believer's *present* perspective. This point was made to me by Wayne Davis, and also by one of the referees.
- ¹⁶ The "ought" which is transgressed in cases of moral negligence is a moral ought, while that transgressed in cases of epistemic negligence is an epistemic ought. Of course we may have moral obligations to be epistemically responsible.
- ¹⁷ Carl Ginet: 1975, *Knowledge, Perception and Memory*, D. Reidel, Dordrecht, p. 28.
- ¹⁸ Ginet (1975, p. 34). This passage, along with the ensuing argument, is discussed by Alston in 'Internalism and Externalism', p. 204ff.
- ¹⁹ Ginet, p. 36.
- ²⁰ Ginet, p. 36.
- ²¹ Ginet, p. 36.
- ²² Ginet, p. 42.
- ²³ H. A. Prichard: 1950, *Knowledge and Perception*, The Clarendon Press, Oxford, p. 86. Prichard is cited by Chisholm in his 1982, 'Knowing That One Knows', in *The Foundations of Knowing*, University of Minnesota Press, Minneapolis, p. 50.
- ²⁴ Roderick Chisholm: 1986, 'The Place of Epistemic Justification', *Philosophical Topics* XIV, 86.
- ²⁵ Chisholm: 1982, 'Knowing That One Knows', p. 55. For another endorsement of level internalism, see Tomis Kapitan: 1985, 'Reliability and Indirect Justification', *Monist* 68.

Alvin Goldman explicitly rejects level internalism in his 1979, 'What is Justified Belief?'. in Moser (ed.), *Empirical Knowledge*. reprinted from G. S. Pappas (ed.), *Justification and Knowledge*, D. Reidel, Dordrecht.

²⁶ The present version is the weakest version of level internalism in the sense that it employs the weakest concepts.

²⁷ A similar argument is made by Van Cleve in his 1984, 'Reliability, Justification, and the Problem of Induction', *Midwest Studies in Philosophy IX*, 563.

²⁸ Chisholm: 1982, 'Knowing That One Knows', p. 55.

²⁹ Roderick Chisholm: 1989, *The Theory of Knowledge*, 3rd edn., Prentice-Hall, Englewood Cliffs, New Jersey, p. 7.

³⁰ Alvin Goldman: 1980, 'The Internalist Conception of Justification', *Midwest Studies in Philosophy V*, 29.

³¹ Goldman, p. 30.

³² Goldman, p. 33.

³³ Pollock: 1986, *Contemporary Theories of Knowledge*, pp. 147–48.

³⁴ This is not to say that old norms, even old initial norms, cannot be replaced by new norms. Thus learning may lead to the adoption of new rules of reasoning which is accompanied by the rejection of old rules. But even this process, if reasonable, will be governed by norms which are already countenanced by the believer. It is an interesting question whether such a process must provide positive justification for a new norm, or whether only a weaker form of coherence with old norms is required.

³⁵ John Pollock argues for this point in his 1986, *Contemporary Theories of Knowledge*, pp. 133–34.

³⁶ I am indebted to David Bennet, Roderick Chisholm, Wayne Davis, Jaegwon Kim, Bindu Madhok, David Martens, Ernest Sosa, Rex Welshon, and Michael Zimmerman for discussion and for their comments on earlier versions of this paper. I am also indebted to the referees of *Synthese*, who provided me with helpful comments on an earlier version.

REFERENCES

- Alston, William: 1985, 'Concepts of Epistemic Justification', *Monist* **68**, 57–89.
- Alston, William: 1986, 'Internalism and Externalism in Epistemology', *Philosophical Topics* **XIV**, 179–221.
- Alston, William: 1988, 'An Internalist Externalism', *Synthese* **74**, 265–83.
- BonJour, Laurence: 1980, 'Externalist Theories of Empirical Knowledge', *Midwest Studies in Philosophy V*, 53–73.
- BonJour, Laurence: 1985, *The Structure of Empirical Knowledge*, Harvard University Press, Cambridge, Massachusetts.
- Chisholm, Roderick: 1977, *Theory of Knowledge*, 2nd edn., Prentice-Hall, Englewood Cliffs, New Jersey.
- Chisholm, Roderick: 1982, *The Foundations of Knowing*, University of Minnesota Press, Minneapolis.
- Chisholm, Roderick: 1986, 'The Place of Epistemic Justification', *Philosophical Topics* **XIV**, 85–92.
- Chisholm, Roderick: 1989, *Theory of Knowledge*, 3rd edn., Prentice-Hall, Englewood Cliffs, New Jersey.

- Foley, Richard: 1984, 'Epistemic Luck and the Purely Epistemic', *American Philosophical Quarterly* **21**, 113–24.
- Ginet, Carl: 1975, *Knowledge, Perception and Memory*, D. Reidel, Dordrecht.
- Goldman, Alvin: 1979, 'What is Justified Belief?', in Moser (ed.), *Empirical Knowledge*, Rowman and Littlefield, Totowa, New Jersey.
- Goldman, Alvin: 1980, 'The Internalist Conception of Justification', *Midwest Studies in Philosophy* **V**, 27–51.
- Kant, Immanuel: 1964, in H. J. Patton (trans.), *Groundwork of the Metaphysics of Morals*, Harper and Row, New York.
- Kapitan, Tomis: 1985, 'Reliability and Indirect Justification', *Monist* **8**, 277–87.
- Luper-Foy, Steven: 1988, 'The Knower, Inside and Out', *Synthese* **74**, 349–67.
- Moser, Paul K., (ed.): 1986, *Empirical Knowledge*, Rowman and Littlefield, Totowa, New Jersey.
- Nagel, Thomas: 1982, 'Moral Luck', in Gary Watson (ed.), *Free Will*, Oxford University Press, Oxford.
- Pappas, G. S. (ed.): 1979, *Justification and Knowledge*, D. Reidel, Dordrecht.
- Pollock, John: 1986, *Contemporary Theories of Knowledge*, Rowman and Littlefield, Totowa, New Jersey.
- Prichard, H. A.: 1950, *Knowledge and Perception*, Clarendon Press, Oxford.
- Van Cleve, James: 1984, 'Reliability, Justification, and the Problem of Induction', *Midwest Studies in Philosophy* **IX**, 555–67.
- Watson, Gary (ed.): 1982, *Free Will*, Oxford University Press, Oxford.
- Zimmerman, Michael J.: 1987, 'Luck and Moral Responsibility', *Ethics* **97**, 374–86.
- Zimmerman, Michael J.: 1988, *An Essay on Moral Responsibility*, Rowman and Littlefield, Totowa, New Jersey.

Department of Philosophy
 Fordham University
 Bronx, NY 10458
 U.S.A.