

MYLAN ENGEL, JR.

PERSONAL AND DOXASTIC JUSTIFICATION
IN EPISTEMOLOGY

(Received 22 October, 1991)

Most epistemologists agree that epistemic justification is necessary for knowledge. This justification requirement is usually formulated in one of two ways:

- (JR1) S knows that p *only if* S is epistemically justified in believing that p.
- (JR2) S knows that p *only if* S's belief that p is epistemically justified.

(JR1) and (JR2) are generally taken to be synonymous formulations of the justification requirement. Accordingly, which formulation is used is thought to be indicative of nothing more than stylistic preference. The synonymy of (JR1) and (JR2) entails what I call the "equivalency thesis". The equivalency thesis asserts:

- (ET) S is epistemically justified in believing that p *iff* S's belief that p is epistemically justified.

According to the equivalency thesis, there is no difference between S being justified in believing that P and S's *belief* that p being justified.

Most epistemologists tacitly embrace the equivalency thesis, since, as we shall see, they jump back and forth between talking about S being justified and S's belief being justified. However, a few are more explicit in their commitment to (ET). For example, Alston says of epistemic justification:

It applies to beliefs, or alternatively to a cognitive subject's having a belief. I shall speak indifferently of S's *belief* that p being justified and of S's being justified in believing that p¹ (emphasis added).

I contend that the equivalency thesis is false and that embracing it

has led philosophers astray in their epistemological theorizing and criticism. In particular, I contend that the current internalist/externalist controversy is directly traceable to the failure on the part of epistemologists to recognize (ET)'s falsity. My aim in the present paper is to defend these contentions.

I

I have suggested that most epistemologists accept the equivalency thesis and its corollary the synonymy of (JR1) and (JR2). Even so, which version of the justification condition they adopt varies, and it varies in a somewhat systematic way. Generally, internalists adopt the (JR1) formulation of the justification requirement. For example, in *Knowledge*, which spells out an internalist coherence theory of justification, Lehrer formulates the justification condition in the following way:

If S knows that p, then S is completely justified in believing that p² (emphasis added).

In *Knowledge and Justification* Pollock, also an internalist, speaks indifferently of beliefs being justified and of person's being justified in beliefs, but when he formally presents conditions, they are generally in the style of (JR1), e.g.:

If P is a prima facie reason for S to believe that Q, and S justifiably believes-that-P and believes-that-Q on the basis of his belief-that-P, then S is justified in believing-that-Q iff he does not believe any defeaters for this prima facie reason³ (emphasis added).

In contrast, externalists seem to embrace the (JR2) formulation. The clearest example of this derives from Goldman's "What Is Justified Belief?", where he asserts: "A theory of justified belief will be a set of principles that specify truth-conditions for the schema |S's belief in p at time t is justified|."⁴

I am not suggesting that internalists use (JR1) exclusively and externalists use (JR2) exclusively, since most epistemologists flipflop indiscriminately between talk of justified persons and talk of justified beliefs.⁵ What I am suggesting is that (JR1) captures something central to internalism and (JR2) captures something central to externalism. Now, if the equivalency thesis is false, then it follows that (JR1) and (JR2) are not synonymous. If (JR1) and (JR2) are not synonymous,

then if I am correct, internalists and externalists have not been analyzing the same concept, and so, rather than disagreeing with each other, they have simply been arguing past each other. Of course, this latter contention rests on the falsity of the equivalency thesis. So, in the next section, I will provide an intuitive argument for (ET)'s falsity.

II

Recall that the equivalency thesis asserts:

(ET) S is epistemically justified in believing that p iff S's belief that p is epistemically justified.

When I look at both sides of this bi-conditional, they seem to be so remarkably different that it is surprising that anyone has taken them to be extensionally equivalent, much less synonymous. Since they have different domains of evaluation,⁶ they do not even purport to be about the same thing. The left-hand side of (ET) is evaluating S, the would-be knower, as being epistemically justified, and hence, it has persons or, more broadly, cognizers as its domain of evaluation. In contrast, the right-hand side of (ET) is evaluating the justificatory status of S's belief, and consequently, it has beliefs as its domain of evaluation.

Given that the two sides of (ET) have different domains of evaluation, a natural question arises as to why they have generally been taken to be extensionally equivalent. The only answer forthcoming is that epistemologists must have thought that a believer cannot be evaluated apart from that which she believes and that a belief cannot be evaluated differently than the cognizer who holds the belief. It seems obvious that both contentions are mistaken. For one, if we consider standard epistemological practice, we find that people are frequently evaluated in terms of the reasoning which leads to their beliefs, rather than the beliefs on which they actually settle. For example, a person is often thought to be justified if she has reasoned well (or if she has done her best to reason well), regardless of what belief she adopts. Regarding the flipside of this epistemological coin, we hear such things as: "Though her belief is a reasonable one to hold, she came to hold it in an epistemically irresponsible way, and consequently, she is unjustified in believing what she does". An example may be helpful here: Consider

Sally the misinformed logic student. Sally is told by her incompetent logic instructor that *modus ponens* is an invalid argument form. Not realizing her instructor's incompetence, she comes to regard *modus ponens* as an invalid form. One day, in a moment of wanton logical abandon, she comes to hold a belief on the basis of a *modus ponens* argument with obviously true premises. Here we have a situation where Sally's belief is perfectly reasonable (since it follows from obviously true premises), but Sally is unjustified in believing it (because, given her situation, she is being epistemically irresponsible in using *modus ponens*). Thus, her belief is evaluated positively, while she is evaluated negatively, in direct contradiction to the claim that beliefs and believers cannot be evaluated independently.

We may summarize the intuitive argument for (ET)'s falsity as follows: Since the two sides of (ET) have different domains of evaluation, they don't even purport to be about the same thing. Therefore, intuitively, they spell out different requirements. Surely, we can agree that the burden of proof lies with the person who wants to butt heads with the intuitive by maintaining that (ET) is true. But the only proof (if you can call it that) that has been offered for (ET)'s truth is the contention that beliefs and believers are not subject to independent epistemic evaluations, and we have just seen with the Sally example that this contention is false. In light of the intuitive evidence for (ET)'s falsity and the lack of any compelling evidence for its truth, I submit that (ET) is false and that, therefore, (JR1) and (JR2) embody inherently different conceptions of epistemic justification. Since (JR1) is concerned with evaluating persons, I will call the kind of epistemic justification associated with it "personal justification". And since (JR2) is evaluative of beliefs, I will call the kind of epistemic justification underlying it "doxastic justification". These two kinds of justification may be characterized as follows:

- (PJ) Personal justification is a normative notion in terms of which *persons* are evaluated from the epistemic point of view.
- (DJ) Doxastic justification is a normative notion in terms of which *beliefs* are evaluated from the epistemic point of view.

In the remainder of this paper, I will attempt to demonstrate the

legitimacy of the personal/doxastic justification distinction, which should, in turn, bolster my claim regarding (ET)'s falsity. Towards this end, I will devote the next section to further clarifying these two kinds of epistemic justification. Then, in section IV, I shall illustrate how distinguishing these two species of justification enables us to sort through and squelch much of the current confusion which surrounds epistemic justification. In the process, it will become clear that internalists have been concerned with personal justification, while externalists have focused on doxastic justification, and that this is what accounts for the radical divergence of their views.

III

Recall that according to (DJ), doxastic justification is a normative notion in terms of which *beliefs* are to be evaluated from the epistemic point of view. So stated, (DJ) does no more than identify doxastic justification as the kind of justification which attaches to beliefs. Of course, for doxastic justification to be a usable notion, we need at least some idea of when a belief possesses it, and so, we need to address the question, "When is a belief doxastically justified?" A trivial, though not wholly uninformative, answer is: A belief is doxastically justified just in case it has positive epistemic status. This suggests that to get a less trivial answer to our question, we need to recast it as follows: "When, from the epistemic point of view, should a belief be evaluated positively?" Since the epistemic point of view is defined by the goal of maximizing truth and minimizing falsity in a large body of beliefs, one might think that the answer to this latter question simply is:

- (A1) S's belief B has positive epistemic status iff B is true.

After all, from the epistemic viewpoint, true beliefs are better than false ones. But an answer like (A1) is essentially nothing more than a restatement of the epistemic goal itself. What we want is a way of evaluating beliefs apart from their actual truth-value, that will indicate which beliefs promote the attainment our dual-pronged goal of gaining truth and avoiding error.

Objective probabilities provide us with such a means of evaluating beliefs.^{7,8} From the epistemic viewpoint, beliefs that are objectively

probable are better than beliefs that are objectively improbable. Thus, a natural answer to our question concerning when a belief has positive epistemic status is:

(A2) S's belief B has positive epistemic status *iff* B has a sufficiently high objective probability of being true.

(A2), together with our earlier trivial observation that a belief is doxastically justified just in case it has positive epistemic status, entails the following nontrivial result:

(DJ') S's belief B is doxastically justified *iff* B has a sufficiently high objective probability of being true.^{9, 10}

(DJ') seems intuitively correct. It captures what, in my opinion, is *the* central idea behind *doxastic* justification by correctly describing which *beliefs* we would want to count as justified from the epistemic standpoint. After all, when we say a *belief* is justified, we are appraising that belief positively as being one that is good to hold for the purposes of gaining truth and avoiding error. According to (DJ'), objectively probable beliefs are good ones to hold for the purposes of gaining truth and avoiding error. This seems right, since objectively probable beliefs are, intuitively, the sort of beliefs best suited for maximizing truth and minimizing falsity within our doxastic corpus. So, *prima facie* (DJ') provides us with a working conception of doxastic justification. Let us now turn to the topic of personal justification.

The expression 'Person S is personally justified in believing that p' is ambiguous. On one reading, it implies that S does, in fact, believe that p and is personally justified in doing so. On a second reading, it asserts that S does not believe that p, but that (given her present cognitive state) she would (or at least could) be personally justified in believing that p, were she to do so.¹¹ For example, we might want to say that S is personally justified in believing that p — in the latter sense — when S has adequate evidence for the belief that p, but has not yet come to believe that p. In order to keep these two senses of S's being personally justified in believing that p separate, I shall adopt the following stipulative terminology: Let the locution 'S is *personally justified* in believing that p' be used exclusively for the former sense that gives existential import to the belief in question, and let the locution 'S is *ex ante*

personally justified in believing that p' be used to capture the second sense where S does not yet hold the belief that p. Throughout this paper when discussing the justificatory status of persons, I will solely be concerned with personal justification, since, unlike *ex ante* personal justification, it has a role to play in the theory of knowledge.¹²

According to our initial characterization (PJ), personal justification is a normative notion in terms of which persons are evaluated from the epistemic viewpoint. In order to flesh out (PJ) and thereby get a hold on the concept of personal justification, we need to reflect on what it is we are doing when we make personal justification evaluations. Let us again start with the obvious. When we evaluate person S as being personally justified in believing that p, we are evaluating S *positively* from the epistemic viewpoint, and when we evaluate S as being personally unjustified in believing that p, we are evaluating S *negatively* from the same viewpoint. Of course, in evaluating S positively from the epistemic viewpoint, we are, in effect, *praising* S epistemically. Similarly, in making the negative evaluation that S is personally unjustified in believing that p, we are *blaming* S epistemically for believing that p. Consequently, at least as a first approximation, we can define personal justification in terms of epistemic praiseworthiness and blameworthiness as follows:

(PJ_i) S is personally justified in believing that p *iff* S is worthy of epistemic praise for believing that p.

(PJ_u) S is personally unjustified in believing that p *iff* S deserves epistemic blame for believing that p.

How, then, do we decide whether a person merits epistemic praise (or blame) for believing a given proposition? We do so on the basis of whether or not she has been epistemically responsible in coming to believe that proposition. If a person comes to believe that p in an epistemically responsible manner (e.g. reasoning carefully, checking her work, considering defeaters, weighing the evidence), she is worthy of epistemic praise and is, therefore, personally justified in believing that p. If, on the other hand, a person comes to believe that p in an epistemically irresponsible manner (e.g. wrecklessly adopting beliefs, ignoring counter-evidence, trusting politicians), she is epistemically blameworthy

and is, therefore, personally unjustified in believing that *p*. Thus, a person's personal justificatory status is a function of whether or not she has proceeded in an epistemically responsible manner when coming to hold a given belief. Accordingly, we can revise (PJ_j) and (PJ_v) as follows:

- (PJ_j) S is personally justified in believing that *p* iff S has come to believe that *p* in an epistemically responsible fashion.
- (PJ_v) S is personally unjustified in believing that *p* iff S has been epistemically irresponsible in coming to believe that *p*.¹³

While there may be other conceptions of personal justification, they are of no interest to me here. The type of personal justification which concerns me in the present paper is that cashed out by (PJ_j) and (PJ_v), since personal justification so-conceived — that brand of justification intimately connected with our notions of epistemic praise, blame, responsibility, and irresponsibility — is, as we shall see, the kind of justification which has served as the impetus for internalism.

It should be noted that the evaluations “personally justified” and “personally unjustified” do not exhaust the class of personal justification evaluations, since very often we hold beliefs for which we deserve neither praise nor blame. Perceptual beliefs (at least most of the time) are a case in point. Typically, our perceptual beliefs come to us as a result of automatic belief-forming cognitive processes over which we have little control. We certainly are not doing anything epistemically reprehensible in holding such beliefs, and so, we are free from epistemic blame for holding them. On the other hand, since we generally do not reflect on our perceptual beliefs or carry out any kind of reasoning procedure to ensure that they are true, we do not deserve epistemic praise for holding them either. Therefore, properly speaking, we are neither justified nor unjustified in holding most of our perceptual beliefs, but rather are *personally ajustified* in holding them. The failure to recognize this third kind of personal justification evaluation has led philosophers to create theories of justification which count persons as being justified, when in fact they are not.

Now that we have a clearer idea of just what personal and doxastic justification are supposed to be, we are in a position to assess the

legitimacy of the personal/doxastic justification distinction. So far I have only offered intuitive considerations in support of the distinction. However, the surest test of a distinction's genuineness is not its intuitiveness, but rather the work that it does. In the remainder of this paper, we shall see that the personal/doxastic justification distinction scores especially well on this latter test.

IV

As previously mentioned, I contend that most epistemologists have conflated personal and doxastic justification and that their doing so has created a number of disputes concerning the nature of epistemic justification, most notably the internalist/externalist dispute. One place where this conflation is particularly noticeable is in certain internalistic counterexamples to the externalist theory called “Process Reliabilism”.

Consider a simplified version of process reliabilism:

- (PR) S's *belief* that *p* is justified in *W* iff it results from a belief-forming cognitive process [BCP] which is *W*-reliable.¹⁴

When is a BCP *W*-reliable? We can answer this question by appealing to the notion of indefinite probability. Unlike definite probabilities which attach to particular propositions, indefinite probabilities are dyadic relations relating classes by specifying the probability of a member of one class being a member of another class. On standard accounts, the indefinite probability of an *A* being a *B* is identified with the actual relative frequency with which *A*'s are *B*'s.¹⁵ Given the foregoing understanding of indefinite probability, we can define a *W*-reliable BCP as follows: A BCP is *W*-reliable iff the indefinite probability of beliefs produced by that BCP in *W* being true beliefs in *W* is high.¹⁶

A proposal like (PR) is epistemologically attractive since: (1) It promises to evince a reduction of the epistemic to the non-epistemic, (2) It sidesteps many of the skeptical problems which so-plague purely internalistic accounts of justification, and (3) It provides a much sought-after conceptual connection between justification and truth. That (PR) provides such a truth connection can be demonstrated as follows: Since (PR) asserts that S's *belief* that *p* is justified in *W* iff it is produced by a

W-reliable BCP, and since, by definition, the indefinite probability of beliefs produced by W-reliable BCPs being true beliefs in W is high, it follows that justified beliefs in W have a high indefinite probability of being true beliefs in W.

Despite its initial attractiveness, (PR) has come under heavy fire from counterexamples who object that (PR)'s *analysans* is neither necessary nor sufficient for epistemic justification. In order to provide a counterexample to (PR)'s necessity, one needs to present a case where intuitively a belief is justified, even though it was produced by an unreliable BCP. The standard objection to (PR)'s necessity runs as follows:

Consider a possible world W where, unbeknownst to us, Descartes' evil demon hypothesis is true and where we are unfortunate enough to reside. In W, virtually all of our beliefs turn out to be false, owing, of course, to the clever manipulations of the demon. Consequently, the BCPs (e.g. perception, memory, and inference) which have produced our beliefs are unreliable in W. So, according to (PR), virtually all of our beliefs are unjustified in W, since they are produced by BCPs that are unreliable in W.

Lehrer and Cohen maintain that such a result is untenable because:

The truth of the demon hypothesis also entails that our experiences and our reasonings are just what they would be if our cognitive processes were reliable, and, therefore, that we would be just as well justified in believing what we do if the demon hypothesis were true as if it were false.¹⁷

Their point is worth belaboring. Our experiences and reasonings in W are, *ex hypothesi*, phenomenologically indistinguishable from the experiences and reasonings we would have in a veridical world W* where we would indeed be justified in holding the beliefs we do. But since our justification for our beliefs, viz. our experiences and reasonings, is exactly the same in both worlds, intuitively we are just as well justified in holding the beliefs we do in W as we are in W*. Since we are justified in holding our beliefs in W, despite the fact that they have all been produced by unreliable BCPs, it follows that reliable production is not a necessary condition for epistemic justification, and therefore, (PR) is false.

Certainly on the face of it, this looks like a devastating counterexample to (PR). In a moment I will argue that this counterexample is, in fact, unsuccessful since it breaks down once the personal/doxastic justification distinction is brought to bear on it, but first I want to

present a second purported counterexample to (PR)'s necessity which Cohen raises.

Consider another demon world W' whose sole inhabitants are A and B. Now imagine that A is a good reasoner, i.e. reasons in accordance with the canons of inductive inference, whereas B engages in confused reasoning, wishful thinking, reliance on emotional attachment, guesswork, etc.¹⁸

As is the case in worlds such as W', unbeknownst to our two hapless inhabitants, the demon sees to it that BCPs like reasoning in accordance with the canons of inductive inference are just as unreliable as BCPs like wishful thinking and confused reasoning.

Since, according to (PR), reliability is the defining feature of justification-conferring BCPs, it follows that in W' the unreliable BCP of reasoning according to the laws of inductive logic is just as *non-justification-conferring* as the equally unreliable BCPs of wishful thinking and confused reasoning. Accordingly, Cohen rightfully notes:

Since the beliefs of A and B are both produced by unreliable processes (the evil demon sees to this), a reliabilist theory of justification must render identical epistemic appraisals of both sets of beliefs.¹⁹

To wit, a reliabilist theory must maintain that neither A's beliefs nor B's beliefs are justified in W'. But, Cohen asserts,

Plainly, this cannot be correct. A's beliefs are conditioned by the evidence whereas B's beliefs are not. A is a good reasoner whereas B is not. A's beliefs are reasonable whereas B's beliefs are not. There is a fundamental epistemic difference between the beliefs of A and the beliefs of B. But the Reliabilist does not have the theoretical means to display this difference.²⁰

The fundamental epistemic difference between A's beliefs and B's beliefs, to which Cohen alludes, is that A's beliefs are supposedly justified whereas B's beliefs are not. Consequently, since A's unreliably produced beliefs are intuitively justified, reliable production is not a necessary condition for epistemic justification, and hence, (PR) is false.

I contend that neither of these seemingly conclusive refutations of (PR)'s necessity is sound, since they both conflate personal justification with doxastic justification. In the first example, we are to suppose that the demon hypothesis is true in our world W and that consequently all of our beliefs have been produced by BCPs which the demon has rendered unreliable in W. So, according to (PR), none of these unreli-

ably produced beliefs are justified in *W*. Interalists find this result unacceptable. The explanation of why it is unacceptable is that intuitively *we* are just as justified in our beliefs in *W* as *we* would be in a veridical world *W**, since our reasonings and experiences are by hypothesis identical in both worlds. Unfortunately for those who offer this explanation, *our* being just as justified in both worlds is completely irrelevant to the question of whether or not our *beliefs* are justified in these worlds. Granted, I agree with the personal evaluation that *we* are just as justified in believing what we do in *W* as we are in *W**, since in such a demon-manipulated environment, we could hardly be faulted for believing what we do. Nevertheless, I contend that in evaluating our *beliefs* as unjustified in *W*, (PR) provides precisely the right result. For in an evil demon world all of our beliefs are produced by highly unreliable BCPs, and so, the indefinite probability of these beliefs being true beliefs is extremely low (at or approaching zero). It strikes me as antithetical to the entire epistemological enterprise to regard such objectively improbable beliefs as having positive epistemic status, i.e. to regard them as being epistemically justified. Obviously, objectively improbable beliefs run counter to the epistemic goal of maximizing truth and minimizing error, since they virtually ensure error, and there is nothing epistemically positive about beliefs which virtually ensure error. This suggests that intuitively our demon-manipulated, objectively improbable *beliefs* are in fact unjustified. Consequently, in the first example, (PR) yields the right *doxastic* justification evaluations, after all.

The second counterexample fails for much the same reason. In that example, we are asked to compare A the good reasoner and B the wishful thinker, both of whom inhabit demon world *W'*, where the demon sees to it that good reasoning and wishful thinking are equally unreliable. Here (PR) yields the result that both A's beliefs and B's beliefs are equally unjustified in *W*. Cohen maintains that this result "plainly cannot be correct." Look once again at what he says.

Plainly, this cannot be correct. A's *beliefs* are conditioned by the evidence whereas B's *beliefs* are not. A is a good reasoner whereas B is not. A's *beliefs* are reasonable whereas B's *beliefs* are not. There is a fundamental epistemic difference between the beliefs of A and the beliefs of B. But the Reliabilist does not have the theoretical means to display this difference²¹ (emphasis added).

Here we are told that there is a fundamental epistemic difference

between the beliefs of A and the beliefs of B and that, therefore, (PR) which fails to take this difference into account must be mistaken. His implicit argument for there being such an epistemic difference between A's beliefs and B's beliefs is roughly that since A is a good reasoner and B is not, the beliefs of A are justified while the beliefs of B are not. In so arguing, he clearly makes the mistake of conflating personal and doxastic justification. His example is instructive, not because it is an objection to (PR), but because it demonstrates the importance of keeping these two kinds of justification separate.

Moreover, the fundamental epistemic difference uncovered by his example is not the one he thinks it is. Since A is a good reasoner and B is a wishful thinker, there is a definite epistemic difference between person A and person B. By reasoning in accordance with the canons of inductive logic, A is presumably forming his beliefs in an epistemically responsible manner. B, on the other hand, in forming wishful beliefs, is being epistemically irresponsible. Hence, A is personally justified in his beliefs in *W'*, whereas B is personally *unjustified* in his beliefs in *W'*. It is crucial to realize, however, that A's being personally justified in his beliefs does *not* entail that his *beliefs* themselves are doxastically justified, especially since, as I will argue, A's beliefs are not doxastically justified.

In the case at hand, the demon sees to it that both A's beliefs and B's beliefs have been produced by equally unreliable BCPs, and because of this unreliable production, A's beliefs are as objectively improbable as B's beliefs. Accordingly, A's beliefs are no better from the epistemic standpoint than B's beliefs, since they are just as likely to result in error as are B's beliefs. According to (DJ'), since both A's beliefs and B's beliefs are objectively improbable, both A's beliefs and B's beliefs are doxastically unjustified. This is precisely as it should be since objectively improbable beliefs do not promote the goal of maximizing truth and minimizing error and, hence, lack positive epistemic status. Since (PR) entails that both A's and B's unreliably produced *beliefs* are unjustified, as a theory of doxastic justification (PR) once again yields the right result. Consequently, this second example also fails to provide us with a counterexample to (PR) as a theory of *justified belief*.

(PR)'s sufficiency has also been challenged. In order to provide a counterexample to (PR) as a sufficient condition for justified belief, one

must present a case where a reliably produced belief is intuitively unjustified. Bonjour, also an internalist, thinks he has found such a case:

Suppose that Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact, the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.²²

Notice that (PR) entails that Norman's belief that the President is in New York City is completely justified. Bonjour attempts to call this result into question with a series of questions:

Is Norman epistemically justified in believing that the President is in New York City, so that his belief is an instance of knowledge? According to the modified externalist position [for our purposes (PR)], we must apparently say that he is. But is this the right result? Are there not still sufficient grounds for a charge of subjective irrationality to prevent Norman's being epistemically justified?²³

By now you can no doubt anticipate my response. There certainly are sufficient grounds for a charge of subjective irrationality to prevent *Norman* from being epistemically justified. After all, Norman has no evidence as to the President's whereabouts. Moreover, he is completely unaware that he has reliable clairvoyant power. So, from his own internal standpoint, it must surely seem as if his belief about the President's present location simply popped into his head out of thin air. And, obviously, it is epistemically irresponsible of Norman to continue to hold such a spontaneously occurring ungrounded belief. All that this shows, however, is that Norman is personally unjustified in believing that the President is in New York City. But Norman's personal unjustifiedness is irrelevant to the doxastic justificational status of his belief. His belief is the result of a highly reliable BCP, to wit, completely reliable clairvoyance. Such beliefs have an extremely high indefinite probability of being true beliefs and, therefore, are good ones to hold from the epistemic standpoint. Notice, Norman's situation with respect to his faculty of clairvoyance is not unlike a young child's situation with respect to her perceptual faculties. What's more, it is usually agreed that the young child's perceptual beliefs are justified, even though she lacks any rationale for them. So, by parity of reason, we should agree that

Norman's *belief* about the President's whereabouts is justified, while nevertheless maintaining that, relative to the other things he believes, he is being epistemically irresponsible in holding the belief and is, therefore, personally unjustified in doing so. Simply put, the fault lies with Norman, not with his belief, and our respective justificatory evaluations should reflect this fact. Therefore, I submit that if (PR) is viewed as an account of doxastic justification, i.e. as a theory of justified *belief*, it remains unscathed by Bonjour's purported counterexample, as well.

V

To sum up, we have seen that, as a theory of doxastic justification, (PR) is immune to all three of the counterexamples examined. What is surprising is that up until now these counterexamples have generally been taken to be decisive. Even Goldman has taken them to be decisive and has, since, abandoned the version of process reliabilism that he formulated in "What Is Justified Belief?". That these counterexamples have been so readily accepted, even by the counterexamples, is, no doubt, a manifestation of the tendency on the part of epistemologists to conflate personal justification with doxastic justification. In each case, the counterexample starts with an internalist, personalist evaluation — *we* are just as justified in *W* as in *W**; *A* is justified in *W'*; *Norman* is unjustified in his belief about the President — and concludes by assuming that the beliefs in question are subject to exactly the same justificatory evaluations as the persons holding them. We now know that such an assumption is both unwarranted and clearly mistaken.

VI

Having seen the importance of distinguishing personal and doxastic justification, one question remains: Which kind of justification is necessary for knowledge? I contend that only doxastic justification is needed for knowledge. Personal justification isn't necessary, because knowledge does not require epistemically praiseworthy behavior on the part of the knower. Even so, personal justification has an important, albeit negative, role to play in the theory of knowledge, which can be seen in the following account of non-Gettierized knowledge:

- (K) S knows that p (in non-Gettier situations) iff:
- (1) p is true,
 - (2) S believes that p,
 - (3) S's belief that p is doxastically justified, and
 - (4) S is *not* personally unjustified in believing that p.

In short, (K) equates knowledge (in non-Gettier situations) with reliably produced true belief which is not irresponsibly held. While a thoroughgoing defense of (K) must await another occasion, we can at least use Norman and Sally as test cases. As Bonjour rightly notes, Norman's subjective irrationality prevents him from knowing that the President is in New York city. (K) yields the same result. According to (K), Norman's reliably produced true belief that the President is in New York City falls short of knowledge, because it is epistemically irresponsible of Norman to hold such a spontaneously occurring ungrounded belief.²⁴ Likewise, Sally fails to know whereof she believes, because, given her situation, it is epistemically irresponsible of her to use *modus ponens*. Once again (K) yields the intuitively correct result. The upshot is that (K) looks quite promising as an analysis of knowledge in non-Gettier situations.

Let us conclude. We have already seen that acknowledging the personal/doxastic justification distinction helps to resolve the internalist/externalist controversy, since it explains why internalists and externalists have such strong competing intuitions regarding epistemic justification, namely, they have been wielding different conceptions of such justification. And, if I am right, the personal/doxastic justification distinction also provides us with the theoretical machinery needed to properly analyze knowledge in non-Gettier situations. With these points in mind, I submit that if we want to make progress in epistemology, we need to keep an eye toward the personal/doxastic justification distinction and tailor our theorizing and criticism accordingly.²⁵

NOTES

¹ William Alston, "Concepts of Epistemic Justification", *The Monist*, Vol. 68, No. 1 (January, 1985), p. 58.

² Keith Lehrer, *Knowledge*, (Oxford: Oxford University Press, 1974), p. 13.

³ John Pollock, *Knowledge and Justification* (Princeton, NJ: Princeton University Press, 1974), p. 44.

⁴ Alvin Goldman, "What is Justified Belief?", *Justification and Knowledge*, ed. George Pappas (Dordrecht: D. Reidel Publishing Company, 1979), p. 3.

⁵ For a case in point see the quote from Stewart Cohen in section IV of the present paper. For a noteworthy exception, see Kent Bach's "A Rationale for Reliabilism", *The Monist*, Vol. 68, No. 2 (April, 1985). Bach recognizes the distinction between justified persons and justified beliefs, but his characterizations of the two kinds of justification differ from mine.

⁶ In *Epistemology and Cognition* Goldman, an externalist, suggests a variety of domains which are subject to epistemic evaluation. The epistemically evaluable domains that he recognizes are: beliefs, methods, psychological processes, hypothesis-forming processes, concept-forming processes, search processes, second-order processes, speech acts, institutional arrangements, and social structures and processes. The one domain he overlooks is: persons. This supports my contention that externalists have been concerned with evaluating beliefs, not persons, while internalists have focused on evaluating persons, not beliefs, and that, thus, unbeknownst to them internalists and externalists have not been discussing the same subject. [See *Epistemology and Cognition*, (Cambridge, Mass. and London: Harvard University Press, 1986), p. 21.]

⁷ Despite their usefulness in other contexts, subjective probabilities are of no use here. After all, in an evil demon world, all of our subjectively probable contingent beliefs are false, and so, subjective probabilities are not conceptually connected to truth. Since subjective probabilities can fail to be truth conducive, they do not provide an effective means of maximizing truth and minimizing falsity.

⁸ Objective probabilities, also known as physical probabilities, depend on objective features of the world and are independent of knowledge and belief. Although various accounts of objective probability have been advanced, I will presuppose a frequency account in the present article. For an edifying discussion of physical probability, see Chapter 4 of John Pollock's *Contemporary Theories of Knowledge* (Totowa, NJ: Rowman and Littlefield, 1986).

⁹ In advancing (DJ'), which identifies doxastically justified belief with objectively probable belief, I must inevitably address the question of "Objectively probable relative to what?", since for a belief to be objectively probable, it must be objectively probable relative to something. I will address this question in Section IV.

¹⁰ In "A Rationale for Reliabilism", *op. cit.*, Bach characterizes what I call "doxastic justification", i.e. the kind of justification which attaches to beliefs, as "whatever property a true, ungettiered belief must possess to qualify as knowledge." In section VI we will see that this characterization leads to counterintuitive results. See footnote 24.

¹¹ Both Goldman and Pollock have noted roughly this same sort of ambiguity, though not in the context of personal justification *per se*. Goldman distinguishes between two uses of the term 'justified', an *ex post* use and an *ex ante* use, which correspond respectively to the two readings I distinguish above. [See his "What is Justified Belief?", *op. cit.*, p. 21.] Pollock distinguishes between justified belief and justifiable belief. "A justifiable belief," he tells us, "is one the believer could become justified in believing if he just put together what he already believes in the right way." [from his *Contemporary Theories of Knowledge*, *op. cit.*, p. 90.]. Accordingly 'justifiable belief' applies only when either (1) S does not yet hold the belief or (2) the belief, which S does hold, is unjustified, because of the way she came to hold it. 'Justified belief' applies only when S does hold the belief. Thus, "justified belief" is like my first reading in that it entails the belief's existence. However, "justifiable belief" differs slightly from my second reading in that "justifiable belief" does not entail the nonexistence of the belief in question.

¹² *Ex Ante* personal justification cannot play a role in an account of knowledge because it entails nonbelief. In entailing nonbelief, it *ipso facto* entails that a necessary condition for knowledge, viz. belief, is not satisfied. Personal justification's role in the theory of knowledge will be discussed in section VI.

¹³ (PJ_i) and (PJ₀) are not being offered as an analysis of the concept "personal justification". Rather, they are intended to point out what the concept of personal justification is. An analysis of personal justification would have to spell out necessary and sufficient conditions for epistemically responsible behavior.

¹⁴ The term 'W-reliable' is shorthand for the more cumbersome 'reliable in W'.

¹⁵ So, e.g., if 9 out of 10 A's are B's, the indefinite probability of an A being a B is 0.9. It should be noted that indefinite probabilities are objective probabilities. To see why, suppose that, in the case at hand, we mistakenly think that only 5 out of 10 A's are B's. In such a situation, we would quite naturally take the probability of an A being a B to be 0.5, but we would be wrong. The probability of an A being a B would remain 0.9, despite what we believe.

¹⁶ This definition of W-reliable BCPs immediately gives rise to the following question: How high must this indefinite probability be for the BCP in question to be W-reliable? Obviously, it must be greater than 0.5. We needn't speculate as to exactly how much greater than 0.5 this indefinite probability needs to be, since none of the arguments in this paper depend on any specific probability threshold or even on there being such a threshold. This is as it should be, for it may turn out that there simply is no context-independent probability necessary for W-reliability, since the degree of justification, and *ipso facto* the degree of BCP reliability, may prove to be context dependent. For an interesting defense of the indexicality of knowledge attributions, see Stewart Cohen's "Knowledge and Context", *The Journal of Philosophy*, Vol. 84 (1986) or his "Knowledge, Context, and Social Standards", *Synthese* Vol. 73 (1987).

¹⁷ Keith Lehrer and Stewart Cohen, "Justification, Trust and Coherence", *Synthese*, 55 (1983), p. 192.

¹⁸ Paraphrased from Stewart Cohen's doctoral dissertation *Justification and Truth*, in manuscript, p. 10.

¹⁹ *Ibid.*, pp. 10–11.

²⁰ *Ibid.*, p. 11.

²¹ Stewart Cohen, *Justification and Truth*, *op. cit.*, p. 11.

²² Laurence Bonjour, "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy*, Vol. 5 (1980), p. 62.

²³ *Ibid.*

²⁴ Bach's characterization of doxastic justification (see footnote 10) precludes personal justification from having a role to play in the theory of knowledge. As a result, on Bach's view, Norman does know that the President is in New York City. Bach embraces this extremely counterintuitive result. See his "A Rationale for Reliabilism", *op. cit.*, pp. 252–253.

²⁵ Earlier versions of this paper were read at Inter-University Centre, Dubrovnik, Yugoslavia and at Karl-Franzens-Universität, Graz, Austria. This paper is much improved as a result of the discussions which ensued at both universities. I am especially indebted to Alfred Schramm whose detailed and penetrating criticisms helped shape the final draft.

Department of Philosophy
Northern Illinois University
Dekalb, IL 60115
 USA