CONCEPTS OF EPISTEMIC JUSTIFICATION

I

Justification, or at least 'justification', bulks large in recent epistemology. The view that knowledge consists of true-justified-belief (+ . . .) has been prominent in this century, and the justification of belief has attracted considerable attention in its own right. But it is usually not at all clear just what an epistemologist means by 'justified', just what concept the term is used to express. An enormous amount of energy has gone into the attempt to specify conditions under which beliefs of one or another sort are justified; but relatively little has been done to explain *what it is* for a belief to be justified, what that is for which conditions are being sought.[1] The most common procedure has been to proceed on the basis of a number of (supposedly) obvious cases of justified belief, without pausing to determine what property it is of which these cases are instances. Now even if there were some single determinate concept that all these theoriests have implicitly in mind, this procedure would be less than wholly satisfactory. For in the absence of an explicit account of the concept being applied, we lack the most fundamental basis for deciding between supposed intuitions and for evaluating proposed conditions of justification. And in any event, as philosophers we do not seek merely to speak the truth, but also to gain an explicit, reflective understanding of the matters with which we deal. We want to know not only when our beliefs are justified, but also what it is to enjoy that status. True, not every fundamental concept can be explicated, but we shall find that much can be done with this one.

And since, as we shall see in this paper, there are several distinct concepts that are plausibly termed "concepts of epistemic justification", the need for analysis is even greater. By simply using 'justified' in an unexamined, intuitive fashion the epistemologist is covering up differences that make important differences to the shape of a theory of justification. We cannot fully understand the stresses and strains in thought about justification until we uncover the most crucial differences between concepts of epistemic justification.

Not all contemporary theorists of justification fall under these strictures. Some have undertaken to give an account of the concept of justification they are using.[2] But none of them provide a map of this entire conceptual territory.

In this paper I am going to elaborate and interrelate several distinct concepts of epistemic justification, bringing out some crucial issues involved in choosing between them. I shall give reasons for disqualifying some of the contenders, and I shall explain my choice of a winner. Finally I shall vouchsafe a glimpse of the enterprise for which this paper is a propadeutic, that of showing how the differences between these concepts make a difference in what it takes for the justification of belief, and other fundamental issues in epistemology.

Before launching this enterprise we must clear out of the way a confusion between one's *being* justified in believing that *p*, and one's *justifying* one's belief that *p*, where the latter involves one's *doing* something to show that *p*, or to show that one's belief was justified, or to exhibit one's justification. The first side of this distinction, on the other hand, is a state or condition one is in, not anything one does or any upshot thereof. I might *be* justified in believing that there is milk on the table because I see it there, even though I have done nothing to show that there is milk on the table or to show that I am justified in believing there to be. It is amazing how often these matters are confused in the literature. We will be concentrating on the "be justified" side of this distinction, since that is of more fundamental epistemological interest. If epistemic justification were restricted to those cases in which the subject carries out a "justification" it would *obviously* not be a necessary condition of knowledge or even of being in a strong position to acquire knowledge. Most cases of perceptual knowledge, for example, involve no such activity.[3]

II

Let's begin our exploration of this stretch of conceptual territory by listing a few basic features of the concept that would seem to be common ground.

(1) 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 *p*. This is the common philosophical concept of belief, in which S's believing that *p* entails neither that S knows that *p* nor that S does not know that *p*. It is not restricted to conscious or occurent beliefs.

(2) It is an evaluative concept, in a broad sense in which this is contrasted with "factual." To say tht S is justified in believing that *p* is to imply that there is something all right, satisfactory, in accord with the way things should be, about the fact that S believes that *p*. It is to accord S's believing a positive evaluative status.

(3) It has to do with a specifically *epistemic* dimension of evaluation. Beliefs can be evaluated in different ways. One may be more or less prudent, fortunate, or faithful in holding a certain belief. Epistemic justification is different from all that. Epistemic evaluation is undertaken from what we might call the "epistemic point of view." That point of view is defined by the aim at maximizing truth and minimizing falsity in a large body of beliefs. The qualification "in a large body of beliefs" is needed because otherwise one could best achieve the aim by restricting one's belief to those that are obviously true. That is a rough formulation. How large a body of beliefs should we aim at? Is any body of beliefs of a given size, with the same truth-falsity ratio, equally desirable, or is it more important, epistemically, to form beliefs on some matters than others? And what relative weights should be assigned to the two aims at maximizing truth and minimizing falsity? We can't go into all that here; in any event, however these issues are settled it remains true that our central cognitive aim is to amass a large body of beliefs with a favorable truth-falsity ratio. For a belief to be epistemically justified is for it, somehow, to be awarded high marks relative to that aim.

(4) It is a matter of degree. One can be more or less justified in believing that p. If, e.g., what justifies one is some evidence one has, one will be more or less justified depending on the amount and strength of the evidence. However in this paper I shall, for the sake of simplicity, treat justification as absolute. You may, if you like, think of this as the degree of justification required for some standard of acceptability.

III

Since any concept of epistemic justification is a concept of some condition that is desirable or commendable from the standpoint of the aim at maximizing truth and minimizing falsity, in distinguishing different concepts of justification we will be distinguishing different ways in which conditions can be desirable from this standpoint. As I see it, the major divide in this terrain has to do with whether believing, and refraining from believing, are subject to obligation, duty, and the like. If they are, we can think of the favorable evaluative status of a certain belief as consisting in the fact that in holding that belief one has fulfilled one's obligations, or refrained from violating one's obligations, to achieve the fundamental aim in question. If they are not so subject, the favorable status will have to be thought of in some other way.

I shall first explore concepts of the first sort, which I shall term 'deontological',[4] since they have to do with how one stands in believing that p,

vis-a-vis duties or obligations. Most epistemologists who have attempted to explicate justification have set out a concept of this sort.[5] It is natural to set out a deontological concept on the model of the justification of behavior. Something I *did* was justified just in case it was *not in violation* of any relevant duties, obligations, rules, or regulations, and hence was not something for which I could rightfully be blamed. To say that my expenditures on the trip were justified is not to say that I was obliged to make those expenditures (e.g., for taxis), but only that it was all right for me to do so, that in doing so I was not in violation of any relevant rules or regulations. And to say that I was justified in making that decision on my own, without consulting the executive committee, is not to say that I was required to do it on my own (though that *may* also be true); it is only to say that the departmental by-laws permit the chairman to use his own discretion in matters of this kind. Similarly, to say that a belief was deontologically justified is not to say that the subject was obligated to believe this, but only that he was permitted to do so, that believing this did not involve any violation of relevant obligations. To say that I am justified in believing that salt is composed of sodium and chlorine, since I have been assured of this by an expert, is not to say that I am obligated to believe this, though this might also be true. It is to say that I am permitted to believe it, that believing it would not be a violation of any relevant obligation, e.g., the obligation to refrain from believing that p in the absence of adequate reasons for doing so. As Carl Ginet puts it, "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."[6]

Since we are concerned specifically with the *epistemic* justification of belief, the concept in which we are interested is not that of *not violating obligations of any sort in believing*, but rather the more specific concept of *not violating "epistemic," "cognitive," or "intellectual" obligations in believing*. Where are such obligations to be found? If we follow out our earlier specification of the "epistemic point of view," we will think of our basic epistemic obligation as that of doing what we can to achieve the aim at maximizing truth and minimizing falsity within a large body of beliefs. There will then be numerous more specific obligations that owe their status to the fact that fulfilling them will tend to the achievement of that central aim. Such obligations might include *to refrain from believing that p in the absence of sufficient evidence* and *to accept whatever one sees to be clearly implied by something one already believes (or, perhaps, is already justified in believing)*.[7] Of course other positions might be taken on this point.[8] One might suppose that there are a number of ultimate, irreducible intellectual duties that cannot be derived from any basic goal of our cognitive life. Or

alternative versions of the central aim might be proposed. Here we shall think in terms of the basic aim we have specified, with more specific obligations derived from that.

Against this background we can set out our first concept of epistemic justification as follows, using 'd' for 'deontological':

I. S is J_d in believing that *p iff* in believing that *p* S is not violating any epistemic obligations.

There are important distinctions between what we may call "modes" of obligation, justification, and other normative statuses. These distinctions are by no means confined to the epistemic realm. Let's introduce them in connection with moral norms for behavior. Begin with a statement of obligation in "objective" terms, a statement of the objective state of affairs I might be said to be obliged to bring about. For example, it is my obligation as a host to make *my guest, G, feel welcome.* Call that underlined state of affairs, 'A'. We may think of this as an *objective* conception of my obligation as a host. I have fulfilled that obligation *iff* G feels welcome.[9] But suppose I did what I sincerely believed would bring about A? In that case surely no one could blame me for dereliction of duty. That suggests a more *subjective* conception of my obligation as *doing what I believed was likely to bring about A.*[10] But perhaps I should not be let off so easily as that. "You should have realized that what you did was not calculated to make G feel welcome." This retort suggests a somewhat more stringent formulation of my obligation than the very permissive subjective conception just specified. It suggests that I can't fulfill my obligation by doing just anything I happen to believe will bring about A. I am not off the hook unless *I did what the facts available to me indicate will have a good chance of leading to A.* This is still a subjective conception in that what it takes to fulfill my obligation is specified from my point of view; but it takes my point of view to range over not all my beliefs, but only my justified beliefs. This we might call a *cognitive* conception of my obligation.[11] Finally, suppose that I did what I had adequate reason to suppose would produce A, and I did produce A, but I didn't do it for that reason. I was just amusing myself, and I would have done what I did even if I had known it would not make G feel welcome. In that case I might be faulted for moral irresponsibility, however well I rate in the other modes. This suggests what we may call a motivational conception of my obligation as *doing what I believed (or was justified in believing) would bring about A, in order to bring about A.*

We may sum up these distinctions as follows:

II. S has fulfilled his *objective* obligation *iff* S has brought about A.

III. S has fulfilled his *subjective* obligation *iff* S has done what he

believed to be most likely to bring about A.

IV. S has fulfilled his *cognitive* obligation *iff* S did what he was justified in believing to be most likely to bring about A.

V. S has fulfilled his *motivational* obligation *iff* S has done what he did because he supposed it would be most likely to bring about A.

We can make analogous distinctions with respect to the justification of behavior or belief, construed as the absence of any violation of obligations.[12] Let's indicate how this works out for the justification of belief.

VI. S is *objectively* justified in believing that *p iff* S is not violating any objective obligation in believing that *p*

VII. *S is subjectively* justified in believing that *p iff* S is not violating any subjective obligation in believing that *p*.

VIII. S is *cognitively* justified in believing that *p iff* S is not violating any congitive obligation in believing that *p*.

IX. S is *motivationally* justified in believing that *p iff* S is not violating any motivational obligation in believing that *p*.

If we assume that only one intellectual obligation is relevant to the belief in question, viz., the obligation to believe that *p* only if one has adequate evidence for *p*, we can be a bit more concrete about this.

X. S is objectively justified in believing that *p iff* S has adequate evidence for *p*.[13]

XI. S is subjectively justified in believing that *p iff* S believes that he possesses adequate evidence for *p*.

XII. S is cognitively justified in believing that *p iff* S is justified in believing that he possesses adequate evidence for *p*.[14]

XIII. S is motivationally justified in believing that *p iff* S believes that *p* on the basis of adequate evidence, or, alternatively, on the basis of what he believed, or was justified in believing, was adequate evidence.

I believe that we can safely neglect XI. To explain why I will need to make explicit what it is to have adeqate evidence for *p*. First a proposition, *q*, *is* adequate evidence for *p* provided they are related in such a way that if *q* is true then *p* is at least probably true. But I *have* that evidence only if I believe that *q*. Furthermore I don't "have" it in such a way as to thereby render my belief that *p* justified unless I know or am justified in believing that *q*. An unjustified belief that *q* wouldn't do it. If I believe that Begin has told the cabinet that he will resign, but only because I credited an unsubstantiated rumour, then even if Begin's having told the cabinet that he

would resign is an adequate indication that he will resign, I will not thereby be justified in believing that he will resign.

Now I might very well *believe* that I have adequate evidence for *q* even though one or more of these conditions is not satisfied. This is an especially live possibility with respect to the first and third conditions. I might mistakenly believe that my evidence is adequate support, and I might mistakenly suppose that I am justified in accepting it. But, as we have just seen, if I am not justified in accepting the evidence for *p* then my believing it cannot render me justified in believing that *p*, however adequate that evidence. I would also hold, though this is perhaps more controversial, that if the evidence is not in fact adequate my having that evidence cannot justify me in believing that *p*. Thus, since my believing that I have adequate evidence is compatible with these non-justifying states of affairs, we cannot take subjective justification, as defined in XI, to constitute epistemic justification.

That leaves us with three contenders. Here I will confine myself to pointing out that there is a strong tendency for J_d to be used in a cognitive rather than a purely objective form. J_d is, most centrally, a concept of freedom from blameworthiness, a concept of being "in the clear" so far as one's intellectual obligations are concerned. But even if I don't have adequate evidence for *p*, I could hardly be blamed for believing that *p* (even assuming, as we are in this discussion, that there is something wrong with believing in the absence of adequate evidence), provided I am justified in supposing that I have adequate evidence. So long as that condition holds I have done the right thing, or refrained from dong the wrong thing, so far as I am able to tell; and what more could be required of me? But this means that it is XII, rather than X, that brings out what it takes for freedom from blame, and so brings out what it takes for being J_d.[15]

What about the motivational form? We can have J_d in any of the first three forms with or without the motivational form. I can have adequate evidence for *p*, and believe that *p*, (XI) whether or not my belief is based on that evidence; and so for the other two. But the motivational mode is parasitic on the other modes, in that the precise form taken by the motivational mode depends on the status of the (supposed) evidence on which the belief is based. This "unsaturated" character of the motivational mode is reflected in the threefold alternative that appears in our formulation of XIII. If S bases his belief that *p* on actually possessed adequate evidence, then XIII combines with X. If the evidence on which it is based is only believed to be adequate evidence, or only justifiably believed to be adequate evidence, then XIII combines with XI or XII. Of course, it may be based on actually possessed adequate evidence, which is justifiably believed to be

such; in which case S is justified in all four modes. Thus the remaining question concerning J_d is whether a "motivational rider" should be put on XII. Is it enough for J_d that S be justified in believing that he has adequate evidence for *p*, or should it also be required that S's belief that *p* be based on that evidence? We will address this question in section V in the form it assumes for a quite different concept of justification.[16]

IV

We have explained *being J_d in believing that p* as *not violating any intellectual obligations in believing that p*. And, in parallel fashion, being J_d in refraining from believing that *p* would consist in not having violated any intellectual obligations in so doing. But if it is possible for me to violate an obligation in refraining from believing that *p*, it must be that I can be obliged, under certain conditions, to believe that *p*. And, by the same token, if I can violate obligations in believing that *p* then I can be obliged to refrain from believing that *p*. And this is the way we have been thinking of it. Our example of an intellectual obligation has been the obligation to refrain from believing that *p* in the absence of adequate evidence. On the other side, we might think of a person as being obliged to believe that *p* if confronted with conclusive evidence that *p* (where that includes the absence of sufficient overriding evidence to the contrary).

Now it certainly looks as if I can be obliged to believe or to refrain from believing, only if this is in my direct voluntary control; only if I can, here and now, believe that *p* or no just by willing (deciding, choosing . . .). And that is the way many epistemologists seem to construe the matter. At least many formulations are most naturally interpreted in this way. Think back, e.g., on Chisholm's formulation of our intellectual obligation (1977, p. 14), cited in n16. Chisholm evisages a person thinking of a certain proposition as a candidate for belief, considering what grounds there might be for belief or refraining from belief, and then effectively choosing belief or absention on the basis of those considerations.[17] Let's call the version of J_d that presupposes direct voluntary control over belief (and thus thinks of an obligation to believe as an obligation to bring about belief here and now), 'J_{dv}' ('v' for 'voluntary').

I find this assumption of direct voluntary control over belief quite unrealistic. There are strong reasons for doubting that belief is usually, or perhaps ever, under direct voluntary control. First, think of the beliefs I acquire about myself and the world about me through experience—through perception, self-consciousness, testimony, and simple reasoning based on these data. When I see a car coming down the street I am not capable of

believing or disbelieving this at will. In such familiar situations the belief-acquisition mechanism is isolated from the direct influence of the will and under the control of more purely cognitive factors.

Partisans of a voluntary control thesis will counter by calling attention to cases in which things don't appear to be so cut and dried: cases of radical underdetermination by evidence, as when a general has to dispose his forces in the absence of sufficient information about the position of enemy forces; or cases of the acceptance of a religious or philosophical position where there seem to be a number of equally viable alternatives. In such cases it can appear that one makes a decision as to what to believe and what not to believe. My view on these matters is that insofar as something is chosen voluntarily it is something other than a belief or abstention from belief. The general chooses to proceed on the working assumption that the enemy forces are disposed in such-and-such a way. The religious convert to whom it is not clear that the beliefs are correct has chosen to live a certain kind of life, or to selectively subject himself to certain influences. And so on. But even if I am mistaken about these kinds of cases, it is clear that for the vast majority of beliefs nothing like direct voluntary control is involved. And so J_{dv} could not possibly be a generally applicable concept of epistemic justification.

If I am right in rejecting the view that belief is, in general or ever, under direct voluntary control, are we foreclosed from construing epistemic justification as freedom from blameworthiness? Not necessarily. We aren't even prevented from construing epistemic justification as the absence of obligation-violations. We *will* have to avoid thinking of the relevant obligations as obligations to believe or refrain from believing, on the model of obligations to answer a question or to open a door, or to do anything else over which we have immediate voluntary control.[18] If we are to continue to think of intellectual obligations as having to do with believing it will have to be more on the model of the way in which obligations bear on various other conditions over which one lacks direct voluntary control but which one can influence by voluntary actions, such conditions as being overweight, being irritable, being in poor health, or having friends. I can't institute, nullify, or alter any of those conditions here and now just by deciding to do so. But I can do things at will that will influence those conditions; and in that way they may be to some extent under my indirect control. One might speak of my being obliged to be in good health or to have a good disposition, meaning that I am obliged to do what I can (or as much as could reasonably be expected of me) to institute and preserve those states of affairs. However since I think it less misleading to say exactly what I mean, I will not speak of our being obliged to weigh a certain amount or to have a good disposition,

or to believe a proposition; I will rather speak of our having obligations to do what we can, or as much as can reasonably be expected of us, to influence those conditions.[19]

The things we can do to affect our believings can be divided into (1) activities that bring influences to bear, or withhold influences from, a particular situation, and (2) activities that affect our belief forming habits. (1) includes such activities as checking to see whether I have considered all the relevant evidence, getting a second opinion, searching my memory for analogous cases, and looking into the question of whether there is anything markedly abnormal about my current perceptual situation. (2) includes training myself to be more critical of gossip, talking myself into being either more or less subservient to authority, and practicing greater sensitivity to the condition of other people. Moreover it is plausible to think of these belief-influencing activities as being subject to intellectual obligations. We might, e.g., think of ourselves as being under an obligation to do what we can (or what could reasonably be expected of us) to make our belief-forming processes as reliable as possible.

All this suggests that we might frame a deontological conception of being epistemically justified in believing that *p*, in the sense that one's believing that *p* is not the result of one's failure to fulfill one's intellectual obligations vis-a-vis one's belief forming and maintaining activities. It would, again, be like the way in which one is or isn't to blame for other conditions that are not under direct voluntary control but which one can influence by one's voluntary activities. I am to blame for being overweight (being irritable, being in poor health, being without friends) only if that condition is in some way due to my own past failures to do what I should to limit my intake or to exercise or whatever. If I would still be overweight even if I had done everything I could and should have done about it, then I can hardly be blamed for it. Similarly, we may say that I am subject to reproach for believing that *p*, provided that I am to blame for being in that doxastic condition, in the sense that there are things I could and should have done, such that if I had done them I would not now be believing that *p*. If that is the case I am unjustified in that belief. And if it is *not* the case, if there are no unfulfilled obligations the fulfilling of which would have inhibited that belief formation, then I am justified in the belief.

Thus we have arrived at a deontological concept of epistemic justification that does not require belief to be under direct voluntary control. We may label this concept 'J_{di}' ('i' for 'involuntary'). It may be more formally defined as follows:

XIV. S is J_{di} in believing that *p* at *t iff* there are no intellectual obligations that (1) have to do with the kind of belief-forming or sus-

taining habit the activation of which resulted in S's believing that p at t, or with the particular process of belief formation or sustenance that was involved in S's believing that p at t, and (2) which are such that:

A. S had those obligations prior to t.
B. S did not fulfill those obligations.
C. If S had fulfilled those obligations, S would not have believed that p at t.[20]

As it stands, this account will brand too many beliefs as unjustified, just because it is too undiscriminating in the counter-factual condition, C. There are ways in which the non-fulfillment of intellectual obligations can contribute to a belief acquisition without rendering the belief unjustified. Suppose that I fail to carry out my obligation to spend a certain period in training myself to observe things more carefully. I use the time thus freed up to take a walk around the neighborhood. In the course of this stroll I see two dogs fighting, thereby acquiring the belief that they are fighting. There was a relevant intellectual obligation I didn't fulfill, which is such that if I had fulfilled it I wouldn't have acquired that belief. But if that is a perfectly normal perceptual belief, it is surely not thereby rendered unjustified.

Here the dereliction of duty contributed to belief-formation simply by facilitating access to the data. That's not the kind of contribution we had in mind. The sorts of cases we were thinking of were those most directly suggested by the two sorts of intellectual obligations we distinguished: (a) cases in which the belief was acquired by the activation of a habit that we would not have possessed had we fulfilled our intellectual obligations; (b) cases in which we acquire, or retain, the belief only because we are sheltered from adverse considerations in a way we wouldn't be if we had done what we should have done. Thus we can avoid counter-examples like the above by reformulating C as follows:

C. If S had fulfilled those obligations, then S's belief-forming habits would have changed, or S's access to relevant adverse considerations would have changed, in such a way that S would not have believed that p at t.

But even with this refinement J_{di} does not give us what we expect of epistemic justification. The most serious defect is that it does not hook up in the right way with an adequate, truth-conducive ground. I may have done what could reasonably be expected of me in the management and cultivation of my doxastic life, and still hold a belief on outrageously inadequate grounds. There are several possible sources of such a discrepancy. First there is what we might call "cultural isolation." If I have grown up in an isolated community in which everyone unhesitatingly accepts the traditions

of the tribe as authoritative, then if I have never encountered anything that seems to cast doubt on the traditions and have never thought to question them, I can hardly be blamed for taking them as authoritative. There is nothing I could reasonably be expected to do that would alter that belief-forming tendency. And there is nothing I could be expected to do that would render me more exposed to counter-evidence. (We can suppose that the traditions all have to do with events distant in time and/or space, matters on which I could not be expected to gather evidence on my own.) I am J_{di} in believing these things. And yet the fact that it is the tradition of the tribe that *p* may be a very poor reason for believing that *p*.

Then there is deficiency in cognitive powers. Rather than looking at the extremer forms of this, let's consider a college student who just doesn't have what it takes to follow abstract philosophical reasoning, or exposition for that matter. Having read Bk. IV of Locke's *Essay*, he believes that it is Locke's view that everything is a matter of opinion, that one person's opinion is just as good as another's, and that what is true for me may not be true for you. And it's not just that he didn't work hard enough on this particular point, or on the general abilities involved. There is nothing that he could and should have done such that had he done so, he would have gotten this straight. He is simply incapable of appreciating the distinction between "One's knowledge is restricted to one's own ideas" and "Everything is a matter of opinion." No doubt teachers of philosophy tend to assume too quickly that this description applies to some of their students, but surely there can be such cases; cases in which either no amount of time and effort would enable the student to get straight on the matter, or it would be unreasonable to expect the person to expend that amount of time or effort. And yet we would hardly wish to say that the student is justified in believing what he does about Locke.

Other possible sources of a discrepancy between J_{di} and epistemic justification are poor training that the person lacks the time or resources to overcome, and an incorrigible doxastic incontinence. ("When he talks like that I just can't help believing what he says.") What this spread of cases brings out is that J_{di} is not sufficient for epistemic justification; we may have done the best we can, or at least the best that could reasonably be expected of us, and still be in a very poor epistemic position in believing that *p*; we could, blamelessly, be believing *p* for outrageously bad reasons. Even though J_{di} is the closest we can come to a deontological concept of epistemic justification if belief is not under direct voluntary control, it still does not give us what we are looking for.

V

Thus neither version of J_d is satisfactory. Perhaps it was misguided all along to think of epistemic justification as freedom from blameworthiness. Is there any alternative, given the non-negotiable point that we are looking for a concept of epistemic evaluation? Of course there is. By no means all evaluation, even all evaluation of activities, states, and aspects of human beings, involves the circle of terms that includes 'obligation', 'permission', 'right', 'wrong', and 'blame'. We can evaluate a person's abilities, personal appearance, temperament, or state of health as more or less desirable, favorable, or worthwhile, without taking these to be within the person's direct voluntary control and so subject to obligation in a direct fashion (as with J_{dv}), and without making the evaluation depend on whether the person has done what she should to influence these states (as with J_{di}). Obligation and blame need not come into it at all. This is most obvious when we are dealing with matters that are not even under indirect voluntary control, like one's basic capacities or bodily build. Here when we use positively evaluative terms like 'gifted' or 'superb', we are clearly not saying that the person has done all she could to foster or encourage the condition in question. But even where the condition is at least partly under indirect voluntary control, as with personal appearance or state of health, we need not be thinking in those terms when we take someone to present a pleasing appearance or to be in splendid health. Moreover, we can carry out these evaluations from a certain point of view. We can judge that someone has a fine bodily constitution from an athletic or from an aesthetic point of view; or that someone's manner is a good one from a professional or from a social point of view.

In like fashion one can evaluate S's believing that *p* as a good, favorable, desirable, or appropriate thing, without thinking of it as fulfilling or not violating an obligation, and without making this evaluation depend on whether the person has done what she could to carry out belief-influencing activities. As in the other cases, it could simply be a matter of the possession of certain good-making characteristics. Furthermore believings can be evaluated from various points of view, including the epistemic, which, as we have noted, is defined by the aim at maximizing truth and minimizing falsity. It may be a good thing that S believes that *p* for his peace of mind, or from the standpoint of loyalty to the cause, or as an encouragement to the redoubling of his efforts. But none of this would render it a good thing for S to believe that *p* from the epistemic point of view. To believe that *p* because it gives peace of mind or because it stimulates effort may not be conducive to the attainment of truth and the avoidance of error.

All of this suggests that we can frame a concept of epistemic justification that is "evaluative," in a narrow sense of that term in which it contrasts with 'deontological', with the assessment of conduct in terms of obligation, blame, right, and wrong. Let's specify an "evaluative" sense of epistemic justification as follows:

XV. S is J_e in believing that *p iff* S's believing that *p*, as S does, is a good thing from the epistemic point of view.

This is a way of being commendable from the epistemic point of view that is quite different from the subject's not being to blame for any violation of intellectual obligations.[21] The qualification "as S does" is inserted to make it explicit that in order for S to be J_e in believing that *p* it need not be the case that any believing of *p* by S would be a good thing epistemically, much less any believing of *p* by anyone. It is rather that there are aspects of *this* believing of *p* by S that make it a good thing epistemically. There could conceivably be person-proposition pairs such that any belief in that proposition by that person would be a good thing epistemically; but this would be a limiting case and not typical of our epistemic condition.

Is there anything further to be said about this concept? Of course we should avoid building anything very substantive into the constitution of the concept. After all, it is possible for epistemologists to differ radically as to the conditions under which one or another sort of belief is justified. When this happens they are at least sometimes using the same concept of justification; otherwise they wouldn't be disagreeing over what is required for justification, though they could still disagree over which concept of justification is most fundamental or most useful. Both our versions of J_d are quite neutral in this way. Both leave it completely open as to what intellectual obligations we have, and hence as to what obligations must not be violated if one is to be justified. But while maintaining due regard for the importance of neutrality I believe that we can go beyond XV in fleshing out the concept.

We can get a start on this by considering the following question. If goodness from an epistemic point of view is what we are interested in, why shouldn't we identify justification with truth, at least extensionally? What could be better from that point of view than truth? If the name of the game is the maximization of truth and the minimization of falsity in our beliefs, then plain unvarnished truth is hard to beat. However this consideration has not moved epistemologists to identify justification with truth, or even to take truth as a necessary and sufficient condition for justification. The logical independence of truth and justification is a staple of the epistemological literature. But why whould this be? It is obvious that a

belief might be J_d without being true and vice versa; but what reason is there for taking J_e to be independent of truth?

I think the answer to this has to be in terms of the "internalist" character of justification. When we ask whether S is justified in believing that *p*, we are, as we have repeatedly been insisting, asking a question from the standpoint of an aim at truth; but we are not asking whether things are in fact as S believes. We are getting at something more "internal" to S's "perspective on the world." This internalist feature of justification made itself felt in our discussion of J_d when we pointed out that to be J_{dv} is to fail to violate any relevant intellectual obligations, *so far as one can tell*, to be J_{dv} in what we call the "cognitive" mode. With respect to J_e the analogous point is that although this is goodness vis-a-vis the aim at truth, it consists not in the beliefs fitting the way the facts actually are, but something more like the belief's being true "so far as the subject can tell from what is available to the subject." In asking whether S is J_e in believing that *p* we are asking whether the truth of *p* is strongly indicated by what S has to go on; whether, given what S had to go on, it is at least quite likely that *p* is true. We want to know whether S had *adequate* grounds for believing that *p*, where *adequate* grounds are those sufficiently indicative to the truth of *p*.

If we are to make the notion of *adequate grounds* central for J_e we must say more about it. A belief has a certain ground, G, when it is "based on" G. What is it for a belief, B, to be *based on* G? That is a difficult question. So far as I know, there is no fully satisfactory general account in the literature, nor am I able to supply one. But we are not wholly at a loss. We do have a variety of paradigm cases; the difficulty concerns just how to generalize from them and just where to draw the line. When one infers *p* from *q* and *thereby* comes to accept *p*, this is a clear case of basing one belief on another. Again, when I come to believe that that is a tree because this visually appears to me to be the case, that is another paradigm; here my belief that that is a tree is based on my visual experience, or, if you prefer, on certain aspects of that experience. The main difficulties arise with respect to cases in which no conscious inference takes place but in which we are still inclined to say that one belief is based on another. Consider, e.g., my forming the belief that you are angry on seeing you look and act in a certain way. I perform no conscious inference from a proposition about your demeanor and behavior to a proposition about your emotional state. Nevertheless it seems plausible to hold that I did learn about your demeanor and behavior through seeing it, and that the beliefs I thereby formed played a crucial role in my coming to believe that you are angry. More specifically it seems that the former beliefs gave rise to the latter belief; that if I hadn't acquired the

former I would not have acquired the latter; and, finally, that if I am asked why I suppose that you are angry I would cite the behavior and demeanor as my reason (perhaps only as "the way he looked and acted"). How can we get this kind of case together with the conscious-inference cases into a general account? We might claim that they are all cases of inference, some of them being unconscious. But there are problems as to when we are justified in imputing unconscious inferences. We might take it that what lets in our problem cases is the subject's disposition to cite the one belief(s) as his reason for the other belief; and then make our general condition a disjunction of conscious inference from q and a tendency to cite q as the reason. But then what about subjects (small children and lower animals) that are too unsophisticated to be able to answer questions as to what their reasons are? Can't their beliefs be based on something when no conscious inference is performed? Moreover this disjunctive criterion will not include cases in which a belief is based on an experience, rather than on other beliefs. A third suggestion concerns causality. In all the cases mentioned thus far it is plausible to suppose that the belief that q was among the causes of the belief that p. This suggests that we might try to cut the Gordian knot by boldly identifying "based on" with "caused by," But this runs into the usual difficulties of simple causal theories. Many items enter into the causation of a belief, e.g., various neuro-physiological happenings, that clearly don't qualify as even part of what the belief is based on. To make a causal account work we would have to beef it up into "caused by q in a certain way." And what way is that? Some way that is paradigmatically exemplified by our paradigms? But how to state this way in such a fashion that it applies equally to the non-paradigmatic cases?[22]

In the face of these perplexities our only recourse is to keep a firm hold on our paradigms, and work with a less than ideally determinate concept of a relationship that holds in cases that are "sufficiently like" the paradigms. That will be sufficient to do the job over most of the territory.[23]

Let's return to "grounds." What a belief is based on we may term the ground of the belief. A ground, in a more dispositional sense of the term, is the sort of item on which a belief can be based. We have already cited beliefs and experiences as possible grounds, and these would seem to exhaust the possibilities. Indeed, some epistemologists would find this too generous already, maintaining that beliefs can be based only on other beliefs. They would treat perceptual cases by holding that the belief that a tree is over there is based on the *belief that* there visually appears to me to be a tree over there, rather than, as we are suggesting, on the visual appearance itself. I can't accept that, largely because I doubt that all perceptual believers have such beliefs about their visual experience,[24] but I can't pause

to argue the point. Suffice it to say that since my opponents' position is, to be as generous as possible, controversial, we do not want to build a position on this issue into the *concept* of epistemic justification. We want to leave open at least the *conceptual* possibility of *direct* or *immediate* justification by experience (and perhaps in other ways also), as well as *indirect* or *mediate* justification by relation to other beliefs (inferentially in the most explicit cases). Finally, to say that a subject *has adequate* grounds for her belief that *p* is is to say that she has other justified beliefs, or experiences, on which the belief could be based and which are strongly indicative of the truth of the belief. The reason for the restriction to *justified* beliefs is that a ground shouldn't be termed adequate unless it can confer justification on the belief it grounds. But we noted earlier that if I infer my belief that *p*, by even impeccable logic, from an *unjustified* belief that *q*, the former belief is not thereby justified.[25]

To return to the main thread of the discussion, we are thinking of S's being J_e in believing that *p* as involving S's having adequate grounds for that belief. That is, we are thinking of the possession of those adequate grounds as constituting the goodness of the belief from the epistemic point of view. The next thing to note is that the various "modes" of J_d apply here as well.

Let's begin by noting an objective-subjective distinction. To be sure, in thinking of J_e as *having truth-indicative grounds within one's "perspective on the world,"* we are already thinking of it as more subjective than flat-out truth. But within that perspectival conception we can set the requirements as more objective or more subjective. There is more than one respect in which the possession of adequate grounds could be "subjectivized". First, there is the distinction between the existence of the ground and its adequacy. S is *objectively* J_e in believing that *p* if S does in fact have grounds that are in fact adequate grounds for that belief. A subjective version would require only that S *believe* one or the other part of this, or both; either (a) that there are (possible) grounds that are in fact adequate and he believes of those grounds that he has them; or (b) that he has grounds that he believes to be adequate; or the combination, (c) that he believes himself to have adequate grounds. Moreover, there are two ways in which the possession-of-grounds belief could go wrong. Confining ourselves to beliefs, one could mistakenly suppose oneself to believe that *p*, or one could mistakenly suppose one's belief that *p* to be justified. Lacking time to go into all these variations, I shall confine this discussion to the subjectivization of adequacy. So our first two modes will be:

XVI. Objective—S does have adequate grounds for believing that *p*.

XVII. Subjective—S has grounds for believing that *p* and he believes them to be adequate.

And here too we have a "justified belief", or "cognitive" variant on the subjective version.

XVIII. Cognitive—S has grounds for believing that *p* and he is justified in believing them to be adequate.

We can dismiss XVII by the same arguments we brought against the subjective version of J_d. The mere fact that I believe, however unjustifiably or irresponsibly, that my grounds for believing that *p* are adequate could scarcely render me justified in believing that *p*. If I believe them to be adequate just because I have an egotistical penchant to overestimate my powers, that could hardly make it rational for me to believe that *p*. But here we will not find the same reason to favor XVIII over XVI. With J_d the cognitive version won out because of what it takes for blameworthiness. But whether one is J_e in believing that *p* has nothing to do with whether he is subject to blame. It depends rather on whether his believing that *p* is a *good thing* from the epistemic point of view. And however justifiably S believes that his grounds are adequate, if they are not then his believing that *p* on those grounds is not a good move in the truth-seeking game. Even if he isn't to blame for making that move it is a bad move nonetheless. Thus J_e is properly construed in the objective mode.

We are also confronted with the question of whether J_e should be construed "motivationally". Since we have already opted for an objective reading, the motivational version will take the following form:

XIX. Motivational—S's belief that *p* is based on adequate grounds.

So our question is whether it is enough for justification that S *have* adequate grounds for his belief, whether used or not, or whether it is also required that the belief be based on those grounds. We cannot settle this question on the grounds that were available for J_{dv}, since with J_e we are not thinking of the subject as being obliged to take relevant consideration into account in *choosing* whether to believe that *p*.

There is something to be said on both sides of this issue. In support of the first, source-irrelevant position (XVI without XIX), it can be pointed out that S's *having a justification* for believing that *p* is independent of whether S does believe that *p*; I can have adequate grounds for believing that *p*, and so *have* a justification, even though I do not in fact believe that *p*. Hence it can hardly be a requirement for having a justification for *p* that my non-existent belief have a certain kind of basis. Likewise my having adequate grounds for believing that *p* is sufficient for this being *a rational thing for me to believe*. But, says the opponent, suppose that S does believe that

p. If simply having adequate grounds were sufficient for this belief to be justified, then, provided S does have the grounds, her belief that *p* would be justified however frivolous the source. But surely a belief that stems from wishful thinking would not be justified, however strong one's (unutilised) grounds for it.[26]

Now the first thing to say about this controversy is that both antagonists win, at least to the extent that each of them is putting forward a viable concept, and one that is actually used in epistemic assessment. There certainly is the concept of *having* adequate grounds for the belief that *p*, whether or not one does believe that *p*, and there equally certainly is the concept of one's belief being based on adequate grounds. Both concepts represent favorable epistemic statuses. *Ceteris paribus*, one is better off believing something for which one has adequate grounds than believing something for which one doesn't. And the same can be said for the contrast between having a belief that is based on adequate grounds and having one that isn't. Hence I will recognize that these are both concepts of epistemic justification, and I will resist the pressure to decide which is *the* concept.

Nevertheless we can seek to determine which concept is more fundamental to epistemology. On this issue it seems clear that the motivational concept is the richer one and thereby embodies a more complete account of a belief's being a good thing from the epistemic point of view. Surely there is something epistemically undesirable about a belief that is generated in an intellectually disreputable way, however adequate the unutilised grounds possessed by the subject. If, possessing excellent reasons for supposing that you are trying to discredit me professionally, I nevertheless believe this, not for those reasons but out of paranoia, in such a way that even if I didn't have those reasons I would have believed this just as firmly, it was undesirable from the point of view of the aim at truth for me to form that belief as I did. So if we are seeking the most inclusive concept of what makes a belief a good thing epistemically, we will want to include a consideration of what the belief is based on. Hence I will take XIX as the favored formulation of what makes a belief a good thing from the epistemic point of view.

I may add that XVI can be seen as derivative from XIX. To simply *have* adequate grounds is to be in such a position that *if* I make use of that position as a basis for believing that *p* I will thereby be justified in that belief. Thus XVI gives us a concept of a potential for XIX; it is a concept of having resources that are sufficient for believing justifiably, leaving open the question of whether those resources are used.

The next point to be noted is that XIX guarantees only *prima facie* justification. As often noted, it is quite possible for my belief that *p* to have

been formed on the basis of evidence that in itself adequately supports *p*, even though the totality of the evidence at my disposal does not. Thus the evidence on which I came to believe that the butler committed the murder might strongly support that hypothesis, but when arriving at that belief I was ignoring other things I know or justifiably believe that tend to exculpate the butler; the total evidence at my disposal is not sufficient support for my belief. In that case we will not want to count my belief as justified all things considered, even though the grounds *on the basis of which* it was formed were themselves adequate. Their adequacy is, so to say, *overriden* by the larger perspectival context in which they are set. Thus XIX gives us *prima facie* justification, what will be justification provided it is not cancelled by further relevant factors. Unqualified justification requires an additional condition to the effect that S does not also have reasons that suffice to override the justification provided by the grounds on which the belief is based. Building that into XIX we get:

XX. Motivational—S's belief that *p* is based on adequate grounds, and S lacks overriding reasons to the contrary.

Even though XX requires us to bring in the unused portions of the perspective, we cannot simplify the condition by ignoring the distinction between what provides the basis and what doesn't, and make the crucial condition something like "The totality of S's perspective provides adequate support." For then we would run up against the considerations that led us to prefer XIX to XVI.

We have distinguished two aspects of our evaluative concept of justification, the strictly evaluative portion—goodness from the epistemic point of view—and the very general statement of the relevant good making characteristic, *based on adequate grounds in the absence of overriding reasons to the contrary*. In taking the concept to include this second component we are opting for the view that this concept, though unmistakably evaluative rather than "purely factual" in character, is not so purely evaluative as to leave completely open the basis on which this evaluative status supervenes. I do not see how to justify this judgment by reference to any more fundamental considerations. It is just that in reflecting on epistemic justification, thought of in evaluative (as contrasted with deontological) terms, it seems clear to me that the range of possible bases for epistemic goodness is not left completely open by the concept, that it is part of what we mean in terming a belief justified, that the belief was based on adequate grounds (or, at least, that the subject had adequate grounds for it).[27] Though this means that J_e is not maximally neutral on the question of what it takes for justification, it is still quite close to that. It still leaves open

whether there is immediate justification and if so on the basis of what, how strong a ground is needed for justification, what dimensions of strength there are for various kinds of grounds, and so on.

Let's codify our evaluative concept of justification as follows:

XXI. S is J_{eg} in believing that *p iff* S's believing that *p*, as S did, was a good thing from the epistemic point of view, in that S's belief that *p* was based on adequate grounds and S lacked sufficient overriding reasons to the contrary.

In the subscript 'g' stands for 'grounds'.

My supposition that all justification of belief involves adequate grounds may be contested. This does seem incontrovertible for beliefs based on other beliefs and for perceptual beliefs based on experience. But what about beliefs in self-evident propositions where the self-evidence is what justifies me in the belief.[28] On considering the proposition that two quantities equal to the same quantity are equal to each other, this seems obviously true to me; and I shall suppose, though this is hardly uncontroversial, that in those circumstances I am justified in believing it. But where are the adequate grounds on which my belief is based? It is not that there are grounds here about whose adequacy we might well have doubts; it is rather that there seems to be nothing identifiable as grounds. There is nothing here that is distinguishable from my belief and the proposition believed, in the way evidence or reasons are distinct from that for which they are evidence or reasons, or in the way my sensory experience is distinct from the beliefs about the physical world that are based on it. Here I simply consider the proposition and straightaway accept it. A similar problem can be raised for normal beliefs about one's own conscious states. What is the ground for a typical belief that one feels sleepy?[29] If one replies "One's conscious of one's feeling of sleepiness," then it may be insisted, with some show of plausibility, that where one is consciously feeling sleepy there is no difference between one's feeling sleepy and one's being conscious that one is feeling sleepy.

This is a very large issue that I will not have time to consider properly. Suffice it to say that one may treat these as limiting cases in which the ground, though real enough, is minimally distinguishable either from the belief it is grounding or from the fact that makes the belief true. In the first person belief about one's own conscious state the ground coincides with the fact that makes the belief true. Since the fact believed is itself an experience of the subject, there need be nothing "between" the subject and the fact that serves as an indication of the latter's presence. The fact "reveals itself" directly. Self-evident propositions require separate treatment. Here I think

that we can take the *way* the proposition appears to one, variously described as "obviously true," "self-evident," and "clear and distinct," as the ground on which the belief is based. I accept the proposition because it *seems* to me so obviously true. This is less distinct from the belief than an inferential or sensory experiential ground, since it has to do with how I am aware of the proposition. Nevertheless there is at least a minimal distinctness. I can form an intelligible conception of someone's failing to believe that p, where p seems obviously true. Perhaps this person has been rendered unduly sceptical by over-exposure to the logical paradoxes.

VI

Let's go back to the idea that the "based on adequate grounds" part of J_{eg} is there because of the "internalist" character of justification. Contrasts between internalism and externalism have been popular in epistemology lately, but the contrast is not always drawn in the same way. There are two popular ways, both of which are distinct from what I have in mind. First there is the idea that justification is internal in that it depends on what support is available for the belief from "within the subject's perspective," in the sense of what the subject knows or justifiably believes about the world.[30] This kind of internalism restricts justification to mediate or discursive justification, justification by reasons. Another version takes "the subject's perspective" to include whatever is "directly accessible" to the subject, accessible just on the basis of reflection; internalism on this version restricts justifiers to what is directly accessible to the subject.[31] This, unlike the first version, does not limit us to mediate justification, since experience can be taken to be at least as directly accessible as beliefs and knowledge.

In contrast to both these ways of drawing the distinction, what I take to be internal about justification is that whether a belief is justified depends on what it is based on (grounds); and grounds must be other psychological state(s) of the same subject. I am not absolutely certain that grounds are confined to beliefs and experiences, even if experiences are not confined to sensations and feelings but also include, e.g., the way a proposition seems obvious to one, and religious and aesthetic experiences; but these are the prime candidates, and any other examples must belong to some kind of which these are the paradigms. So in taking it to be conceptually true that one is justified in believing that p *iff* one's belief that p is based on an adequate ground, I take justification to be "internal" in that it depends on the way in which the belief stems from the believer's psychological states, which are "internal" to the subject in an obvious sense. What would be an externalist contrast with this kind of internalism? We shall see one such contrast

in a moment, in discussing the relation of J_{eg} to reliabilism. Moreover, it contrasts with the idea that one can be justified in a certain belief just because of the status of the proposition believed (necessary, infallible). My sort of internalism is different from the first one mentioned above, in that experiences as well as beliefs can figure as grounds. And it is different from the second if, as I believe, what a belief is based on may not be directly accessible. This will be the case if, as seems plausible, much belief formation goes on below the conscious level. It would seem, e.g., that, as we move about the environment, we are constantly forming short-term perceptual beliefs without any conscious monitoring of this activity.

The most prominent exponents of an explicitly non-deontological conception of epistemic justification have been reliabilists, who have either identified justification with reliability[32] or have taken reliability to be an adequate criterion of justification.[33] The reliability that is in question here is the reliability of belief formation and sustenance.[34] To say that a belief was formed in a reliable way is, roughly, to say that it was formed in a way that can be depended on generally to form true rather than false beliefs, at least from inputs like the present one, and at least in the sorts of circumstances in which we normally find ourselves.[35] Thus if my visual system, when functioning as it is at present in yielding my belief that there is a tree in front of me, generally yields true beliefs about objects that are fairly close to me and directly in front of me, then my present belief that there is a tree in front of me was formed in a reliable manner.

Now it may be supposed that J_{eg}, as we have explained it, is just reliability of belief formation with an evaluative frosting. For where a belief is based on adequate grounds that belief has been formed in a reliable fashion. In fact, it is plausible to take reliability as a *criterion* for adequacy of grounds. If my grounds for believing that *p* are not such that it is generally true that beliefs like that formed on grounds like that are true, they cannot be termed 'adequate'. Why do we think that wanting State to win the game is not an adequate reason for supposing that it has won, whereas the fact that a victory has been reported by several newspapers is an adequate reason? Surely it has something to do with the fact that beliefs like that when formed on the first sort of grounds are not *generally* true, while they are *generally* true when formed on grounds of the second sort. Considerations like this may lead us to suppose that J_{eg}, in effect, identifies justification with reliability.[36]

Nevertheless the internalist character of justification prevents it from being identified with reliability, and even blocks an extensional equivalence. Unlike justification, reliability of belief formation is not limited to cases in which a belief is based on adequate grounds within the subject's

psychological states. A reliable mode of belief formation *may* work through the subject's own knowledge and experience. Indeed, it is plausible to suppose that all of the reliable modes of belief formation available to human beings are of this sort. But it is quite conceivable that there should be others. I might be so constituted that beliefs about the weather tomorrow which apparently just "pop into my mind" out of nowhere are in fact reliably produced by a mechanism of which we know nothing, and which does not involve the belief being based on anything. Here we would have reliably formed beliefs that are not based on adequate grounds from within my perspective, and so are not J_{eg}.

Moreover, even within the sphere of beliefs based on grounds, reliability and justification do not necessarily go together. The possibility of divergence here stems from another feature of justification embodied in our account, the way in which unqualified justification requires not only an adequate ground but also the absence of sufficient overriding reasons. This opens up the possibility of a case in which a belief is formed on the basis of grounds in a way that is in fact highly reliable, even though the subject has strong reasons for supposing the way to be unreliable. These reasons will (or may) override the *prima facie* justification provided by the grounds on which the belief was based. And so S will not be justified in the belief, even though it was reliably generated.

Consider, in this connection, a case presented by Alvin Goldman.[37]

> Suppose that Jones is told on fully reliable authority that a certain class of his memory beliefs are almost all mistaken. His parents fabricate a wholly false story that Jones suffered from amnesia when he was seven but later developed *pseudo*-memories of that period. Though Jones listens to what his parents say and has excellent reasons to trust them, he persists in believing the ostensible memories from his seven-year-old past.

Suppose that Jones, upon recalling his fifth birthday party, believes that he was given an electric train for his fifth birthday because, as it seems to him, he remembers being given it.[38] By hypothesis, his memory mechanism is highly reliable, and so his belief about his fifth birthday was reliably formed. But this belief is not adequately supported by the *totality* of what he justifiably believes. His justifiable belief that he has no real memory of his first seven years overrides the support from his ostensible memory. Thus Jones is not J_{eg} in his memory belief, because the "lack of overriding reasons to the contrary" requirement is not satisfied. But reliability is subject to no such constraint. Just as reliable mechanisms are not restricted to those that work through the subject's perspective, so it is not a requirement on the reliability of belief-formation that the belief be adequately supported by the totality of the subject's persepctive. However many and however

strong the reasons Jones has for distrusting his memory, the fact remains that his memory beliefs are still reliably formed. Here is another way in which the class of beliefs that are J_{eg} and the class of reliably formed beliefs can fail to coincide.[39]

I would suggest that, of our candidates, J_{eg} most fully embodies what we are looking for under the heading of "epistemic justification". (1) Like its deontological competitors it is an evaluative concept, in a broad sense, a concept of a favorable status from an epistemic point of view. (2) Unlike J_{dv} it does not presuppose that belief is under direct voluntary control. (3) Unlike J_{di}, it implies that the believer is in a strong epistemic position in believing that *p*, i.e., that there is something about the way in which he believes that *p* that renders it at least likely that the belief is true. Thus it renders it intelligible that justification is something we should prize from an epistemic point of view. (4) Unlike the concept of a reliable mode of belief formation it represents this "truth-conducivity" as a matter of the belief's being based on an adequate ground within the subject's own cognitive states. Thus it recognizes the "internalist" character of justification; it recognizes that in asking whether a belief is justified we are interested in the prospects for the truth of the belief, given what the subject "has to go on." (5) Thus the concept provides broad guidelines for the specification of conditions of justification, but within those guidelines there is ample room for disagreement over the precise conditions for one or another type of belief. The concept does not leave us totally at a loss as to what to look for. But in adopting J_{eg} we are not building answers to substantive epistemological questions into the concept. As the only candidate to exhibit all these desiderata, J_{eg} is clearly the winner.

VII

It may be useful to bring together the lessons we have learned from this conceptual exploration.

1. Justifying, an activity of showing or establishing something, is much less central for epistemology than is "being justified," as a state or condition.

2. It is central to epistemic justification that *what justifies* is restricted to the subject's "perspective," to the subject's knowledge, justified belief, or experience.

3. Deontological concepts of justification are either saddled with an indefensible assumption of the voluntariness of belief (J_{dv}) or allow for cases in which one believes that *p* without having any adequate ground for the belief (J_{di}).

4. The notion of one's belief being based on adequate grounds incorporates more of what we are looking for in a concept of epistemic justification than the weaker notion of having adequate grounds for belief.

5. Justification is closely related to reliability, but because of the perspectival character noted in 2., they do not completely coincide; much less can they be identified.

6. The notion of believing that *p* in a way that is good from an epistemic point of view in that the belief is based on adequate grounds (J_{eg}) satisfies the chief desiderata for a concept of epistemic justification.

VIII

The ultimate payoff of this conceptual exploration is the increased sophistication it gives us in dealing with substantive epistemological issues. Putting our scheme to work is a very large enterprise, spanning a large part of epistemology. In conclusion I will give one illustration of the ways in which our distinctions can be of help in the trenches. For this purpose I will restrict myself to the broad contrast between J_{dv} and J_{ep}.

First, consider what we might term "higher-level requirements" for S's being justified in believing that *p*. I include under that heading all requirements that S know or justifiably believe something *about* the epistemic status of *p*, or about the strength of S's grounds for *p*. This would include requirements that S be justified in believing that:

1. R is an adequate reason for *p* (where R is alleged to justify S's belief that *p*).[40]
2. Experience *e* is an adequate indication that *p* (where *e* is alleged to justify S's belief that *p*).[41]

On J_{eg} there is no temptation to impose such requirements. If R *is* an adequate reason (*e is* an adequate indication), then if one believes that *p* on that basis, one is *thereby* in a strong position, epistemically; and the further knowledge, or justified belief, that the reason is adequate (the experience is an adequate indication), though no doubt quite important and valuable for other purposes, will do nothing to improve the truth-conduciveness of one's believing that *p*. But on J_{dv} we get a different story. If it's a question of being blamless in believing that *p*, it can be persuasively argued that this requires not only forming the belief on what is in fact an adequate ground, but doing so in the light of the realization that the ground is an adequate one. If I decide to believe that *p* without knowing whether the ground is adequate, am I not subject to blame for proceeding irresponsibly in my doxastic behavior, whatever the actual strength of the ground? If the higher-

level requirements are plausible only if we are using J_{dv}, then the dubiousness of that concept will extend to those requirements.[42]

In the above paragraph we were considering whether S's being justified in believing that his ground is adequate is a *necessary* condition of justification. We can also consider whether it is sufficient. Provided that S is justified in believing that his belief that *p* is based on an adequate ground, G, does it make any difference, for his being justified in believing that *p*, whether the ground *is* adequate? Our two contenders will line up here as they did on the previous issue. For J_{eg} the mere fact that S is justified in supposing that G is adequate will cut no ice. What J_{eg} requires is that S *actually be* in an epistemically favorable position; and although S's being justified in supposing G to be adequate is certainly good evidence for that, it doesn't *constitute* being in such a position. Hence J_{eg} requires that the ground of the belief actually be an adequate one. As for J_{dv}, where it is a question of whether S is blameworthy in believing that *p*, what is decisive is how S's epistemic position appears within S's perspective on the world. If, so far as S could tell, G is an adequate ground, then S is blameless, i.e., J_{dv}, in believing that *p* on G. Nothing else could be required for justification in that sense. If S has chosen his doxastic state by applying the appropriate principles in the light of all his relevant knowledge and justified belief, then he is totally in the clear. Again the superior viability of J_{eg}, as over against J_{dv}, should tip the scales in favor of the more objective requirement of adequacy.[43]

William P. Alston

Syracuse University

NOTES

1. Of late a number of theorists have been driving a wedge between what it is to *be* P or what *property* P is, on the one hand, and what belongs to the *concept* of P or what is the meaning of 'P' on the other. Thus it has been claimed (Kripke, 1972) that *what heat is* is determined by the physical investigation into the nature of heat, whether or not the results of that investigation are embodied in our *concept* of heat or in the meaning of 'heat'. I shall take it that no such distinction is applicable to epistemic justification, that here the only reasonable interpretation to be given to 'what it is' is 'what is involved in the concept' or 'what the term means'. If someone disagrees with this, that need not be a problem. Such a person can simply read 'what concept of justification is being employed' for 'what justification is taken to be'.

2. I think especially of Chisholm (1977), ch. 1; Ginet (1975), Ch. III; Goldman (1979), (1980); Wolterstorff (1983).

3. It may be claimed that the activity concept is fundamental in another way, viz., by virtue of the fact that one is justified in believing that *p* only if one is *capable* of carrying out a justification of the belief. But if that were so we would be justified in far fewer beliefs than we suppose. Most human subjects are quite incapable of carrying out a justification of any perceptual or introspective beliefs.

4. I am indebted to Alvin Plantinga for helping me to see that this term is more suitable than the term 'normative' that I had been using in earlier versions of this paper. The reader should be cautioned that 'deontological' as used here does not carry the contrast with 'teleological' that is commonin ethical theory. According to that distinction a deontological ethical theory, like that of Kant's, does not regard principles of duty or obligation as owing their status to the fact that acting in the way they prescribe tends to realize certain desirable states of affairs. Whereas a teleological theory, like Utilitarianism, holds that this is what renders a principle of obligation acceptable. The fact that, we are not using 'deontological' with this force is shown by the fact that we are thinking of epistemic obligations as owing their validity to the fact that fulfilling them would tend to lead to the realization of a desirable state of affairs, viz., a large body of beliefs with a favorable truth-falsity ratio.

5. See Chisholm (1977), ch. 1; Ginet (1975), ch. 3; Wolterstorff (1983). An extended development of a deontological concept of epistemic justification is to be found in Naylor (1978). In my development of deontological concepts in this paper I have profited from the writing of all these people and from discussions with them.

6. (1975), p. 28. See also Ayer (1956), pp. 31–34; Chisholm (1977), p. 14; Naylor (1978), p. 8.

7. These examples are mant to be illustrative only; they do not necessarily carry the endorsement of the management.

8. Here I am indebted to Alvin Plantinga.

9. A weaker objective conception would be this. My obligation is to do what in fact is *likely* to bring out A. On this weaker conception I could be said to have fulfilled my obligation in (some) cases in which A is not forthcoming.

10. We could also subjectivize the aimed at result, instead of or in addition to subjectivizing what it takes to arrive at that result. In this way one would have subjectively fulfilled one's obligation if one had done what one believed to be one's obligation. Or, to combine the two moves to the subjective, one would have subjectively fulfilled one's obligation if one had done what one believed would lead to the fulfillment of what one believed to be one's obligation. But sufficient unto the day is the distinction thereof.

11. I would call this "epistemic obligation," except that I want to make these same distinctions with respect to epistemic justification, and so I don't want to repeat the generic term for one of the species.

12. Since we are tacitly restricting this to epistemic justification, we will also be, tacitly, restricting ourselves to intellectual obligations.

13. Since this is all on the assumption that S does believe that *p*, we need not add that to the right hand side in order to get a sufficient condition.

14. Note that XI, XII, and some forms of XIII are in terms of higher-level beliefs about one's epistemic status *vis-a-vis p*. There are less sophisticated sorts of subjectivization. For example:

S is subjectively justified in believing that *p iff* S believes that *q*, and *q* is evidence for *p*.

(For the reason this does not count as having adequate evidence see the next paragraph in the text.)

Or even more subjectively:

> S is subjectively justified in believing that *p iff* S believes that *q* and bases his belief that *p* on his belief that *q*.

The definitions presented in the text do not dictate what we should say in the case in which S does not have the higher level belief specified in XI and XII, but satisfies either of the above conditions. A thorough treatment of modes of normative status would have to go into all of this.

15. We have been taking it that to be, e.g., subjectively or cognitively justified in believing that *p* is to not be violating any subjective or cognitive obligations in believing that *p*. That means that if we opt for cognitive justification we are committed to giving a correspondingly cognitive formulation of what intellectual obligations one has. But that isn't the only way to do it. We could leave all the obligations in a purely objective form, and vary the function that goes from obligation to justification. That is, we could say that one is subjectively justified if one believes that one has not violated an (objective) obligation (or, perhaps believes something that is such that, given one's objective obligations, it implies that none of those obligations have been violated). And a similar move could be made for the other modes.

16. Here are a couple of examples of the attraction of XII for J_d. Chisholm (1977) presents an informal explanation of his basic term of epistemic evaluation, 'more reasonable than' in terms of an "intellectual requirement." The explanation runs as follows.

> One way, then, of re-expressing 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*. (14)

The point that is relevant to our present discussion is that Chisholm states our basic intellectual requirement in what I have called "cognitive" rather than "objective" terms; and with a motivational rider.

> 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. (14)

The "requirement" is that one *try one's best* to bring this about, rather than that one do bring it about. I take it that to try my best to bring about a result, R, is to do what, so far as I can tell, will bring about R, insofar as that is within my power. (It might be claimed that so long as I do what I believe will bring about R I am trying my best, however irresponsible the belief. But it seems to me that so long as I am not acting on the best of the indications available to me I am not "trying my best.") The motivational rider comes in too, since unless I do what I do *because* I am taking it to (have a good chance to) lead to R, I am not trying at all to bring about R.

Of course, Chisholm is speaking in terms of fulfilling an intellectual obligation rather than, as we have been doing, in terms of not violating intellectual obligations. But we are faced with the same choice between our "modes" in either case.

For a second example I turn to Wolterstorff (1983). Wolterstorff's initial formulation of a necessary and sufficient condition of justification (or, as he says, "rationality") for an "eluctable" belief of S that *P* is: *S lacks adequate reasons for ceasing from believing that p.* (164). But then by considerations similar to those we have

just adduced, he recognizes that even if S does not in fact have adequate reason for ceasing to believe that *p* he would still be unjustified in continuing to hold the belief if he were "rationally obliged" to believe that he does have adequate reason to cease to believe that *p*. Moreover Wolterstorff recognizes that S would be justified in believing that *p* if, even though he does have adequate reason to cease from believing that *p* he is rationally justified in supposing that he doesn't. Both these qualifications amount to recognizing that what is crucial is not what reasons S has in fact, but what reasons S is justified in supposing himself to have. The final formulation, embodying these and other qualifications, runs as follows:

> A person s is rational in his eluctable and innocently produced belief *Bp* if and only if S does believe *p* and either:
>
> (i) S neither has nor ought to have adequate reason to cease from believing *p*, and is not rationally obliged to believe that he *does* have adequate reason to cease; or
>
> (ii) S does have adequate reason to cease from believing *p* but does not realize that he does, and is rationally justified in that. (168)

17. See also Ginet (1975), p. 36.

18. Note that I am not restricting the category of what is within my immediate voluntary control to "basic actions." Neither of the actions just mentioned would qualify for that title. The category includes both basic actions and actions that involve other conditions, where I can satisfy those other conditions, when I choose, just at the moment of choice. Thus my point about believing is not just that it is not a basic action, but that it is not even a non-basic action that is under my effective immediate control. Whatever is required for my believing that there will never be a nuclear war, it is not something that I can bring about immediately by choosing to do so; though, as I am about to point out, I can affect my believings and abstentions in a more long range fashion.

19. For other accounts of the indirect voluntary control of beliefs see Naylor (1978), pp. 19–20; Wolterstorff (1983), pp. 153–55.

20. Our four "modes" can also be applied to J_{di}. Indeed, the possibilities for variation are even more numerous. For example, with respect to the *subjective* mode we can switch from the objective fact to the subject's belief with respect to (a) the circumstances of a putative violation, (b) whether there was a violation, and (c) whether the violation was causally related to the belief-formation in question. We will leave all this as an exercise for the reader.

21. I must confess that I do not find 'justified' an apt term for a favorable or desirable state or condition, when what makes it desirable is cut loose from considerations of obligation and blame. Nevertheless, since the term is firmly ensconced in the literature as the term to use for any concept that satisfies the four conditions set out in section II, I will stifle my linguistic scruples and employ it for a non-deontological concept.

22. There are also problems as to where to draw the line. What about the unconscious "use" of perceptual cues for the depth of an object in the visual field or for "size constancy"? And however we answer that particular question, just where do we draw the line as we move farther and farther from our initial paradigms?

23. For some recent discussion of 'based on' see Swain (1981), ch. 3, and Pappas (1979). One additional point I do need to make explicit is this. I mean 'based on' to

range over both what initially gave rise to the belief, and what sustains it while it continues to be held. To be precise one should speak of *what the belief is based on at time t*. If *t* is the time of acquisition one is speaking of what gave rise to the belief; if *t* is later than that one is speaking of what sustains it.

24. For an interesting discussion of this point see Quinton (1973), ch. 7. My opponent will be even more hard pressed to make out that beliefs about one's own conscious experience are based on other beliefs. His best move here would be either to deny that there are such beliefs or to deny that they are based on anything.

25. No such restriction would be required just for having grounds (of some sort). Though even here the word 'ground' by itself carries a strong suggestion that what is grounded is, to some extent, supported. We need a term for anything a belief might be based on, however vainly. 'Ground' carries too much positive evaluative force to be ideally suitable for this role.

26. For some recent discussion of this issue see Harman (1973) ch. 2; Lehrer (1974) ch. 6; Firth (1978); Swain (1981) ch. 3; Foley (1984).

27. Even though we have opted for the 'based on' formulation as giving us the more fundamental concept of epistemic justification, we have also recognized the 'has adequate grounds' formulation as giving us a concept of epistemic justification. Either of these will introduce a "basis of evaluative status" component into the concept.

28. This latter qualification is needed, because I might accept a self-evident proposition on authority. In that case I was not, so to say, taking advantage of its self-evidence.

29. We are not speaking here of a belief that one *is* sleepy. There a ground is readily identifiable—one's feeling of sleepiness.

30. See Bonjour (1980), Kornblith (1985), Bach (1985).

31. See Goldman (1980), Chisholm (1977), ch. 4, pp. 63–64; Ginet (1975), 34–37.

32. Swain (1981), ch. 4.

33. Goldman (1979).

34. For simplicity I shall couch the ensuing formulations solely in terms of belief formation, but the qualification 'or sustenance' is to be understood throughout.

35. These two qualifications testify to the difficulty of getting the concept of reliability in satisfactory shape; and there are other problems to be dealt with, e.g., how to identify the general procedure of which the present belief formation is an instance.

36. An alternative to explicating 'adequate' in terms of reliability would be to use the notion of conditional probability. G is an adequate ground for a belief that *p* just in case the probability of *p* on G is high. And since adequacy is closely related both to reliability and to conditional probability, they are presumably closely related to each other. Swain (1981) ch. 4, exploits this connection to explicate reliability in terms of conditional probability, though in a more complex fashion than is indicated by these brief remarks.

37. (1979) p. 18.

38. If you have trouble envisaging his trusting his memory in the face of his parents's story, you may imagine that he is not thinking of that story at the moment he forms the memory belief.

39. In the article in which he introduces this example Goldman modifies the "reliability is a criterion of justification" view so that it will accommodate the example. The modified formulation runs as follows:

> If S's belief in *p* at *t* results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which had it been used by S in addition to the process actually used, would have resulted in S's not believing *p* at *t*, then S's belief in *p* at *t* is justified. (p. 20)

On this revised formulation, being formed by a reliable process is sufficient for justification only if there is no other reliable process that the subject could have used and such that if he had used it he would not have come to believe that *p*. In the case cited there is such a reliable process, viz., taking account of the strong reasons for believing one's memory of pre-seven years old events to be unreliable. The revised reliability criterion yields the correct result in this case. However this move leaves unshaken the point that in this case Jones's belief *is* reliably formed but unjustified. That remains true, whatever is to be said about the revised criterion.

40. See e.g., Armstrong (1973), p 151; Skyrms (1967), p. 374.

41. See, e.g., Sellars (1963), pp. 168–69; Bonjour (1978), pp. 5–6; Lehrer (1974), pp. 103–05.

42. In my paper, "What's Wrong with Immediate Knowledge?" *Synthese*, vol. 55, no. 2 (May, 1983), pp. 73–95, I develop at much greater length this kind of diagnosis of Bonjour's deployment of a higher-level requirement in his argument against immediate knowledge (Bonjour, 1978).

43. Ancestors of this paper were presented at SUNY at Albany, SUNY at Buffalo, Calvin College, Cornell University, University of California at Irvine, Lehigh University, University of Michigan, University of Nebraska, Syracuse University, and the University of Western Ontario. I wish to thank members of the audience in all these institutions for their helpful comments. I would like to express special appreciation to Robert Audi, Carl Ginet, George Mavrodes, Alvin Plantinga, Fred Schmitt, and Nicholas Wolterstorff for their penetrating comments on earlier versions.

REFERENCES

Armstrong, D. M. (1973)—*Belief, Truth, and Knowledge.* London: Cambridge University Press.

Ayer, A. J. (1956)—*The Problem of Knowledge.* London: Macmillan.

Bach, K. (1985)—"A Rationale for Reliabilism," *The Monist*, vol. 68 no. 2 (April, 1985).

Bonjour, L. (1978)—"Can Empirical Knowledge Have a Foundation?", *Amer. Philos. Quart.*, vol. 15, pp. 1–13.

Chisholm, R. M. (1977)—*Theory of Knowledge*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, Inc.

Dretske, F. (1981)—*Knowledge and the Flow of Information.* Cambridge, MA: MIT Press.

Firth, R. (1978)—"Are Epistemic Concepts Reducible to Ethical Concepts?". In A. I. Goldman & J. Kim, eds., *Values and Morals.* Dordrect: D. Reidel.

Foley, R. (1984)—"Epistemic Luck and the Purely Epistemic". *Amer. Philos. Quart.*

Ginet, C. (1975)—*Knowledge, Perception, and Memory*. Dordrecht: D Reidel Publ. Co.

Goldman, A. I. (1979)—"What Is Justified Belief?". In G. S. Pappas. ed., *Justification and Knowledge*, Dordrecht: D. Reidel Pub. Co.

—(1980)—"The Internalist Conception of Justification". *Midwest Studies in Philosophy*, vol. V, pp. 27–51.

Harman, G. (1973)—*Thought*. Princeton: Princeton University Press.

Kornblith, H. (1985)—"Ever Since Descartes," *The Monist*, vol. 68, no. 2 (April, 1985).

Kripke, S. A. (1972)—"Naming and Necessity," in *Semantics of Natural Language*, ed. D. Davidson & G. Harman. Dordrecht: D. Reidel Pub. Co.

Lehrer, K. (1974)—*Knowledge*. New York: Oxford University Press.

Naylor, M. B. (1978)—"Epistemic Justification." Manuscript.

Pappas, G. S. (1979)—"Basing Relations." In G. S. Pappas, ed., *Justification and Knowledge*. Dordrecht: D. Reidel Pub. Co.

Pollock, J. (1974)—*Knowledge and Justification*. Princeton: Princeton University Press.

Quinton, A. (1973)—*The Nature of Things*. London: Routledge & Kegan Paul.

Sellars, W. (1963)—*Science, Perception, and Reality*. London: Routledge & Kegan Paul.

Skyrms, B. (1967)—"The Explication of 'S Knows that P'." *Journ. Philos.*, vol. 64, pp. 373–89.

Swain, M. (1981)—*Reasons and Knowledge*. Ithaca, NY: Cornell University Press.

Wolterstorff, N. (1983)—"Can Belief in God Be Rational If It Has No Foundations?", in *Faith and Rationality*, ed. A. Plantinga & N. Wolterstorff. Notre Dame: University of Notre Dame Press.