Philosophical Perspectives, 2, Epistemology, 1988

POSITIVE EPISTEMIC STATUS AND PROPER FUNCTION

Alvin Plantinga
University of Notre Dame

Knowledge, so we thought for untold generations,[1] is justified true belief; and even in this enlightened post-Gettier era, we still think justification and knowledge intimately related. But what sort of thing is this alleged 'justification'? How shall we understand it? Contemporary epistemologists don't often focus attention on the *nature* of epistemic justification (although they often ask under what conditions a given belief has it); and when they do, they display deplorable diversity. Some claim that justification is a matter of epistemic dutifulness, others that it goes by coherence, and still others that it is conferred by reliability. In what follows I shall argue that none of the above is the correct answer, and suggest an alternative.

But how shall we initially locate this quality or quantity I mean to discuss; how shall we initially pin down epistemic justification? First 'justification', 'justified' and allied terms are terms of epistemic appraisal[2]; to say that a proposition is *justified* for a person is to say that it has what Roderick Chisholm calls 'positive epistemic status' for him; his holding that belief in his circumstances is *right*, or *proper*, or *acceptable*, or *approvable*, or *up to standard*. What we appraise are a person's *beliefs* (more exactly, his *believings*), as well as his skepticisms and (to use another Chisholmian term) his withholdings, his failings to believe. The evidentialist objector to theistic belief,[3] for example, claims that a theist who believes in God without sufficient evidence is so far forth unjustified in that belief; thus he offers a negative appraisal of the belief (when held without

sufficient evidence) or of its holder—claiming, perhaps, that the believer in those circumstances has flouted some duty, or (more charitably) that she is suffering from a sort of neurosis or other cognitive dysfunction. In the same way we may appraise the belief that all contemporary flora and fauna arose by way of random genetic mutation and natural selection; and of course the less spectacular beliefs of everyday life are subject to similar evaluation and appraisal. Although Chisholm doesn't mention it, we also appraise the *degree* to which a person believes a proposition; if I believe that Homer was born before 800 B.C. with the same fervor that I believe that New York City is larger than Cleveland, then (given my epistemic circumstances) my degree of confidence in the former proposition is unjustified. We evaluate a person's beliefs and degrees of belief as warranted, or justified, or rational, or reasonable, contrasting them with beliefs that are unwarranted, unjustified, irrational, or unreasonable.

Secondly, epistemic justification or positive epistemic status comes in degrees. Some of my beliefs have more by way of positive epistemic status for me than others: for example, my belief that I live in Indiana has more by way of positive epistemic status or warrant, for me, than my belief that Homer was born before 800 B.C. (This is not to say that I am not equally *within my rights* in accepting these two beliefs to the *degrees* to which I do in fact accept them; I believe the former much more firmly than the latter.) But then we can distinguish degrees of positive epistemic status, at least for a given person.[4]

And thirdly, among the fundamental concepts of the theory of *knowledge* we find, naturally enough, the concept of knowledge. There is wide (though less than universal) agreement that true belief is necessary but not sufficient for knowledge. But then what more is required? What is this quantity enough of which (Gettier problems, perhaps, aside) epistemizes true belief? (We can't properly assume that it is a *simple* property or quantity; perhaps it is like a vector resultant of one or more others.) Whatever exactly this further element or quantity may be, it is either epistemic justification or something intimately connected with it. So perhaps the natural procedure would be just to *baptize* this element, whatever it is, 'epistemic justification'. But this would be misleading. The term 'justification' suggests duty, obligation, permission, and rights—the whole deontological stable. Furthermore, one of the main contending theories

or pictures here (a theory or picture going back at least to Descartes) explicitly identifies the quantity in question with *aptness for epistemic duty fulfillment*; to use the term 'justification', then, as a name for the quantity in question would be to give this theory a confusing and unwarranted (if merely verbal) initial edge over its rivals. I shall therefore borrow Chisholm's more neutral term "positive epistemic status" as my official name for the quantity in question. Positive epistemic status, then, initially and to a first approximation, is a normative property that comes in degrees, enough of which is what epistemizes true belief. Now contemporary epistemology is dominated by three fundamentally different basic ideas as to what positive epistemic status is; after arguing that each of these three is deeply flawed, I shall go on to propose a more satisfactory alternative.

I. Chisholmian Internalism[5]

Over the past 25 years or so, Roderick Chisholm (clearly the dean of contemporary epistemologists) has presented a series of ever more refined and penetrating accounts of the central notions of the theory of knowledge. His work is surely as good a place to start as any. Chisholm belongs to an *internalist* tradition going back at least to Descartes. There is more than one important internalist tradition; but according to the dominant tradition, the Cartesian tradition, positive epistemic status is essentially connected with the fulfillment of epistemic duty, with the satisfaction of noetic obligation. This tradition is therefore the natural home of the use of term 'justification' for positive epistemic status, for according to this tradition what epistemizes true belief lies in the near neighborhood of epistemic duty fulfillment. Chisholm is an internalist; but the notion of internalism is less than wholly clear, and there is more than one variety of it.[6] But here I speak of Cartesian internalism, a whole system or congeries of ideas. Central to this system, I think, is the idea that (truth aside) whether a person has knowledge is *up to him* and within his control. More exactly, the central idea is that whether or not my beliefs *have positive epistemic* status for me is up to me and within my control. Perhaps I am the victim of a Cartesian demon or a subject in an Alpha Centaurian cognitive experiment, so that my beliefs are for the most part wildly wide of the mark; nevertheless I can

still do my epistemic duty and thereby still achieve a state in which my beliefs have positive epistemic status for me. According to the internalist, we need give no hostages to fortune when it comes to justification or positive epistemic status; here our fate is in our own hands. Being justified in her beliefs (unlike, say, having a sunny disposition) is not something that just *happens* to a person; it is a result of her own efforts.

But if this is so, then something else must also be so—in fact three things must be so. First, our beliefs must be, at least to a large extent, within our voluntary control. The basic internalist idea is that it is up to me whether I adopt beliefs that are justified for me. But then which beliefs I adopt must be up to me; for any or most beliefs that suggest themselves, that come within my purview, it must be within my power to accept them and within my power to reject them. Second, if it is appropriately up to me whether my beliefs are justified, then I must be able to tell, somehow, whether a given belief *is* justified. Justification must be a property such that I can determine, with confidence and from the inside, so to speak, whether or not a belief or candidate for belief has it. The fates may conspire to deceive me; I could be wrong about whether there is an external world, or a past, or whether there are other persons; for all I can tell, I may be the victim of a malevolent Cartesian demon who delights in deception. Justification, however, is a different kettle of fish; whatever my problems with truth, at least I can determine whether my beliefs are justified. Thirdly, consider Chisholmian epistemic principles: non-contingent propositions stating the conditions under which a person's beliefs enjoy (one or another degree of) positive epistemic status with respect to her. Given the basic internalist idea and a couple of plausible assumptions, it follows that I can discover these principles, or at any rate can discover *some* of these principles *a priori*, just by reflection. I need not resort to a *posteriori* investigation (perhaps of a broadly psychological or anthropological sort); I need not know that my faculties are reliable or functioning properly. All three of these aspects of internalism are reflected in Chisholm's work.[7]

Chisholm presents his epistemology by way of stating and commenting upon epistemic principles: non-contingent conditionals whose antecedents specify a relation between a person **S**, a proposition **A**, and certain circumstances in which **S** finds himself, and whose consequents specify that **A** has a certain epistemic status for **S**—

certainty, perhaps, or *acceptability*, or *being evident*, or *being beyond reasonable doubt*. He begins by introducing an undefined technical locution: "**p** is more reasonable than **q** for **S** at **t**". Here the values for **p** and **q** will be *properties*: such properties, for example, as *believing that Albuquerque is in New Mexico* and *withholding the belief that Albuquerque is in New Mexico*—that is, believing neither that proposition nor its denial. (When one believes or withholds something, then, Chisholm assumes[8], there is indeed something one believes or withholds; call such things 'propositions'.) Given the notion of a proposition and 'is more reasonable than' as an undefined locution, Chisholm goes on to define a battery of "terms of epistemic appraisal" as he calls them: 'certain', 'beyond reasonable doubt', 'evident', 'acceptable', and so on. A proposition **A** is beyond reasonable doubt for a person at a time **t**, for example, if and only if it is more reasonable for him to accept that proposition then than to withhold it; **A** has some presumption in its favor for him at **t** just if accepting it then is more reasonable than accepting its negation. The epistemological principles Chisholm presents are formulated by way of these terms of epistemic appraisal; and the whole process culminates in a definition or analysis of knowledge.

Now Chisholm introduces 'is more reasonable than' as an undefined locution; nonetheless, of course, it has a sense, as he uses it, and a sense fairly close to the sense it has in English. The main thing to see is that as Chisholm uses it, this locution pertains to epistemic *duty* or *requirement* or *obligation*. In *Foundations of Knowing* (hereafter FK), his most recent full dress presentation of his epistemology, he says that "Epistemic reasonability could be understood in terms of the general requirement to try to have the largest possible set of logically independent beliefs that is such that the true beliefs outnumber the false beliefs. The principles of epistemic preferability are the principles one should follow if one is to fulfill this requirement." (7). In his earlier *Theory of Knowledge* (2nd edition (1976); hereafter TK) Chisholm puts it as follows: "We may assume," he says,

> that every person is subject to a purely intellectual requirement: that of trying his best to bring it about that for any proposition *p* he considers, he accepts *p* if and only if *p* is true (TK p. 14);

and he adds

> One might say that this is the person's responsibility *qua* intellectual being....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*.'

Reasonability, therefore, is a *normative* concept. More precisely, it is a *deontological* concept; it pertains to requirement, duty, or obligation— *epistemic* duty or obligation. (Of course it doesn't follow that the normative character involved is strictly *moral*; perhaps an epistemic requirement is not a moral duty; perhaps it is a *sui generis* form of obligation.[9]) And Chisholm's fundamental contention here is that a certain requirement, or responsibility, or duty, or obligation lies at the basis of such epistemic notions as evidence, justification, positive epistemic status, and knowledge itself; for of course he analyzes knowledge in terms of positive epistemic status plus truth. To say, for example, that a proposition **p** is *acceptable* for a person at a time is to say that he is so situated, then, that it is not the case that he can better fulfill his epistemic duty by withholding than by accepting **p**; to say that **p** is *beyond reasonable doubt* for him is to say that he is so situated that he can better fulfill his intellectual responsibility by accepting **p** than by withholding **p**.

The suggestions made in FK and TK do not agree as to what our intellectual requirement is; neither, furthermore, is exactly right.[10] The basic idea, however, is that our epistemic duty or requirement is to try to achieve and maintain a certain condition—call it 'epistemic excellence'— which may be hard to specify in detail, but consists fundamentally in standing in an appropriate relation to truth. This is a duty I have "*qua* intellectual being"—that is, just by virtue of being the sort of creature that is capable of grasping and believing (or withholding) propositions. We must pay a price for our exalted status as intellectual beings; with ability comes responsibility. And the idea, presumably, is that *all* intellectual beings have this responsibility: angels, devils, Alpha Centaurians, what have you—all are subject to this requirement or obligation.

According to Chisholm, then, positive epistemic status is a matter of aptness for fulfillment of epistemic duty. A proposition has positive epistemic status for me, in certain circumstances, to the extent that I can fulfill my epistemic duty by accepting it in those circumstances.

This duty or obligation or requirement, furthermore, is one of *trying* to bring about a certain state of affairs. One's duty as an intellectual being is not that of *succeeding* in bringing it about that (say) one has a large set of beliefs, most of which are true; it is instead that of *trying* to bring about this state of affairs. My requirement is not to *succeed* in achieving and maintaining intellectual excellence; my requirement is only to try to do so. Presumably the reason is that it may not be within my power to succeed. Perhaps I don't know how to achieve intellectual excellence; or perhaps I do know how but simply can't do it. So my duty is only to *try* to achieve it.

But how shall I try to achieve epistemic excellence? What, more concretely, must I do? Chisholm's answer: follow the principles of epistemic preferability—the epistemic principles he has repeatedly tried to state. "The principles of epistemic preferability are the principles one should follow if one is to fulfill this requirement [i.e., the requirement to try to achieve epistemic excellence]" (FK p. 7). On several occasions and in several contexts Chisholm suggests that his epistemic principles are instances of the Reidian idea that we ought to trust our epistemic nature unless we have reason not to: "Our perceptual principles are instances of the more general truth: 'it is reasonable to trust the senses until one has positive reason for distrusting them'" (FK p. 23). How are we to understand this? As follows, I think. We human beings find ourselves with a battery of epistemic impulses, inclinations and dispositions. Upon being appeared to redly, for example, I find myself believing that I am appeared to in that way by some object, an object that is red. Upon considering an elementary truth of arithmetic or logic, I find myself believing that it is true; upon being asked what I had for breakfast this morning, I find myself believing that it was eggs on toast. Our natures are such that for each of a wide variety of circumstances there are certain beliefs we are strongly disposed or inclined to form; and when we find ourselves in these circumstances, we find ourselves with those beliefs. Now Chisholm's idea, I think, is that in these circumstances (given that we have no indication that our epistemic nature is flawed and misleading) it is reasonable to trust our epistemic nature—i.e., to believe as nature prompts us to. But of course 'reasonable' is, he says, to be understood in terms of epistemic duty or obligation: accordingly, in the circumstances in question, the dutiful thing to do is to fall in with our natural inclinations and accept the beliefs nature inclines us towards. In a nutshell, then: we have a

duty to try to achieve epistemic excellence; in a wide variety of circumstances that duty can be best fulfilled by accepting the beliefs our nature inclines us to accept; and a belief or candidate for belief has positive epistemic status for me to the extent that in accepting it I can fulfill my epistemic duty.

This is a simple and attractive picture of the nature of justification and positive epistemic status. It is easy to see, however, that it can't be correct. The fundamental idea is that positive epistemic status is a matter of *aptness for the fulfillment of epistemic duty or obligation.* This may be an *element* or *moment* in positive epistemic status; but I think we can easily see that it cannot be the whole story. I shall give three examples to argue the point.[11]

First, suppose **S** knows that nine out of ten Frisians cannot swim and that Feike is a Frisian. He is aware of the fact that he knows these things, and that they disconfirm

(1) Feike can swim.

He has no evidence of any kind for (1); no perceptual evidence, no propositional evidence, no testimonial evidence—no evidence at all. Nevertheless, (1) seems overwhelmingly attractive to him; it seems wholly and obviously true; it has all the phenomenological *panache* of *modus ponens* itself. For **S** very much admires swimming and swimmers; and, due to a psychological malfunction (**S** is himself a very poor swimmer), he has a powerful tendency to assume, of anyone he likes or admires, that she is an excellent swimmer. **S** isn't aware of this malfunction and has no reason to think he suffers from it; and his lack of awareness is in no way due to epistemic malfeasance or lack of epistemic dutifulness. In fact **S** is extremely concerned with his epistemic duty. He is eager to achieve epistemic excellence, to bring it about that he is in the right relation to the truth. He is trying his level best to do so; indeed, he is nearly fanatic on the subject and devotes what many would consider an inordinate share of his energy to trying to achieve epistemic excellence.

Now what, under these conditions, would be the dutiful thing for **S** to do? Obviously, he should accept (1). Is he so situated that he can better fulfill his obligation to try to achieve epistemic excellence by *withholding* than by accepting (1)? Surely not. (l) seems utterly and obviously true to him; and while he knows that he knows some things that disconfirm it, the same holds for all sorts of propositions he knows to be true. (Thus, for example, he also knows that Tietje

is a Frisian and has himself personally witnessed her winning the 100 meter freestyle at the Olympics.) It is true that his cognitive faculties are playing him false here, but he has no inkling of this fact, and his lack of this self-knowledge is in no way due to epistemic carelessness or other dereliction of epistemic duty. The way for him to try to achieve epistemic excellence in these circumstances, surely, is for him to act on what he nonculpably believes about how best to achieve this end. But (l) seems utterly and obviously true; so, naturally enough, he believes that the way to achieve the end in question is to accept (1). Indeed, there may be no other proposition such that **S** can better fulfill his duty to the truth by accepting it than by accepting (l); in that case, on Chisholm's official account of positive epistemic status, (1) would be *certain* for **S**. (1), therefore, is in these circumstances overwhelmingly apt for fulfillment of epistemic duty; if positive epistemic status were what Chisholm says it is, (l) would have positive epistemic status in *excelsis* for **S**. But surely it doesn't. Even if it is true, **S** certainly does not *know* that it is, and the idea that under these conditions it could have certainty, the maximal degree of positive epistemic status for him, seems wholly fantastic.

It is clear, therefore, that (l) has little by way of positive epistemic status for **S**; at any rate it has little or none of the sort of epistemic status enough of which confers knowledge. Still, it does seem to have *some* kind of positive status for him. We could certainly say that S is *permitted* to accept (l); he is violating no duty in accepting it; he is entirely within his epistemic rights in accepting it. He is *justified* in accepting it in the sense that he has a *right* to accept it; he is doing his epistemic best in accepting it, thus fulfilling his duty to the truth. Nevertheless the proposition in question doesn't have the sort of status for him enough of which (together with truth) constitutes knowledge. I don't mean merely that he doesn't know it (even if it turns out to be true); it is rather that the sort of status it has for him is not such that even if it had more or the maximal degree of *that* sort of status, then he *would* know it. To have the status for him required by *knowledge*, something quite different is demanded. We might say that he has **permissive** justification in accepting this proposition—using the term 'permissive' to indicate that the sort of justification he has is such that in accepting the proposition in question he is entirely within his epistemic rights and is flouting no epistemic duty. Nonetheless, the proposition has little by way of positive epistemic status for him.

A second example: according to Chisholm, something is *appearing* to someone, or appearing in a certain way to her, if (very roughly) it is appropriately causing her to be appeared to in that way (FK 16-17). Now suppose that, due to cerebral malfunction or the machinations of a Cartesian evil demon, I have a strong tendency or impulse to believe

(2) Nothing is appearing redly to me

whenever I am appeared to redly. This tendency is even stronger than the tendency normal people display, in the same circumstances, to believe the appropriate denial of (2); it seems to me utterly obvious, under these conditions, that there isn't anything appearing red to me. Furthermore, I have not the faintest inkling of this defect in my nature, and my failure to be aware of it is in no way due to lack of epistemic dutifulness. (Indeed, we may add in this case, as in the last, that doing my epistemic duty is the main passion of my life.) Then, surely, the dutiful thing for me to do, under those circumstances, would be to accept (2). But though I would be permissively justified in accepting that proposition, it would have little by way of positive epistemic status for me; surely it wouldn't have the sort of positive epistemic status, for me, enough of which is sufficient (with truth) for knowledge. The problem is that my cognitive faculties are not working properly. I display cognitive malfunction, so that no matter how magnificently I do my epistemic duty, no matter how hard I try, I won't have much by way of positive epistemic status. Or rather, I will have the justification that goes with doing my best to do my duty; I will be within my rights, not properly subject to blame or censure. I will have permissive justification. Indeed, I will have more than permissive justification; in trying as hard as I did to achieve epistemic excellence, I performed works of epistemic supererogation. But no amount of dutifulness or supererogatory effort is sufficient for the kind of positive epistemic status necessary for knowledge; for that an element of quite another kind is required.

A final example: Paul is so constructed (again, due to brain lesion or demon or mad Alpha Centaurian scientist) that when he is appeared to in one sense modality he forms beliefs appropriate to another. When he is aurally appeared to in the way in which one is appeared to upon hearing church bells, he has a nearly ineluctable tendency to believe that there is something that is appearing to him in that fashion, and that thing is orange—bright orange. He

doesn't know about this defect in his epistemic equipment, and his lack of awareness is in no way due to indolence, or carelessness, or wishful thinking, or any other dereliction of epistemic duty. As a matter of fact, Paul is unusually dutiful, unusually concerned about doing his epistemic duty; fulfilling this duty is the main passion of his life. Add that those around him suffer from a similar epistemic deficiency. They have all been manipulated in this way by demons or Alpha Centaurians; or they live in Alaska and all suffer from similar lesions due to radioactive fallout from a Soviet missile test. Now suppose Paul is appeared to in the church bell fashion and forms the belief that he is being appeared to in that way by something that is orange. Surely, in these conditions, this proposition is such that accepting it is an appropriate way of doing his epistemic duty, of trying to achieve epistemic excellence. Nevertheless that proposition has next to nothing by way of positive epistemic status for him. Paul is beyond reproach; he has done his duty as he saw it; he is permissively justified, and more. Nevertheless there is a kind of quality this belief lacks—a kind crucial for knowledge. For *that* sort of status, it isn't sufficient to satisfy one's duty and do one's epistemic best. Paul can be ever so conscientious about his epistemic duties and still be such that his beliefs do not have positive epistemic status.

Clearly enough, we can vary the above sorts of examples. Perhaps you think that what goes with satisfying duty *in excelsis* is *effort*; perhaps (in a Kantian vein) you think that genuinely dutiful action demands acting contrary to inclination. Very well, alter the above cases accordingly. Suppose, for example, that Paul (due to lesion, demon or Alpha Centaurian) nonculpably believes that his nature is deeply misleading. Like the rest of us, he has an inclination, upon being appeared to redly, to believe that there is something red lurking in the neighborhood; unlike the rest of us, he believes that this natural inclination is misleading and that on those occasions there really isn't anything that is thus appearing to him. He undertakes a strenuous regimen to overcome this inclination; after intense and protracted effort he succeeds: upon being appeared to redly he no longer believes that something red is appearing to him. His devotion to duty costs him dearly. The enormous effort he expends takes its toll upon his health; he is subject to ridicule and disapprobation on the part of his fellows, who view his project as at best Quixotic; his wife protests his unusual behavior and finally leaves him for someone less epistemically nonstandard. Determined to do what is

right, however, Paul heroically persists in doing what he is nonculpably convinced is his duty. It is obvious, I take it, that even though Paul is unusually dutiful in accepting, on a given occasion, the belief that nothing red is appearing to him, that belief has little by way of positive epistemic status for him.

What these examples show, I think, is that positive epistemic status is not or is not merely a matter of aptness for fulfillment of epistemic duty or obligation. That there *are* such duties or obligation seems eminently plausible, although tough problems attend this notion[12] and it isn't easy to say in any detail what our epistemic duties might be. Aptness for the fulfillment of such duties, however, is at most one aspect or moment of positive epistemic status.[13] It is obvious that a proposition can be maximal with respect to aptness for duty fulfillment for me, but nonetheless enjoy little by way of positive epistemic status.

II. Coherentism

A second suggestion as to the nature of positive epistemic status is made by the **coherentist**. There are, of course, many brands of coherentism and many ways to construe coherence;[14] I don't have the space to canvass them here. We can construe the coherentist in either of two ways: (1) she may agree with the Chisholmian that aptness for duty fulfillment is a moment in positive epistemic status, adding that the other element is provided or determined by coherence, or (2) she may hold that positive epistemic status is not complex and that it is determined solely by coherence. What she claims, taken the second way, is that a belief has positive epistemic status, for me, to the extent that it coheres with the appropriate system of beliefs. (She could then hold either that positive epistemic status *just is* coherence, or that it is instead a normative or axiological property that *supervenes* upon coherence.) I shall argue that coherentism is mistaken if taken the first way; but if mistaken that way, then it is also mistaken if taken the other.

An initial and important problem is that there are few serious attempts to *state* coherentism, few serious attempts to say what this alleged coherence relation is. Most coherentists are decently reticent about the nature of coherence; we are typically told that it is more than mere logical consistency but less than mutual entailment;

beyond this most coherentists maintain a decorous silence.[15] I think we can see, however, that on any plausible account of coherence, coherentism is unacceptable as an account either of warrant or of positive epistemic status. According to the weak version of coherentism, coherence is the source of one element or component of positive epistemic status, the other being aptness for epistemic duty fulfillment; on the strong version coherence just is or is the sole source of positive epistemic status. If the strong version is true, the weak follows trivially; I shall therefore argue that the weak version is false. The chief problem for coherentism, as I see it, is that coherence is thought of as a relation that obtains just among *beliefs*. But that means that one of my beliefs could have a great deal of positive epistemic status by virtue of standing in the coherence relation to the right body of beliefs or propositions, *no matter how it was related to my experience*; This, however, is clearly mistaken. Once again I shall proceed by way of examples. It is easy to see, I think, that a given belief could have next to nothing by way of positive epistemic status for a person, even though she was wholly dutiful in holding the belief, and even though the belief in question is coherent with the rest of her beliefs.

Consider, first, someone who suffers from a cognitive dysfunction: whenever he is appeared to redly (to use Chisholm's term), he forms the belief that no one is ever appeared to redly. It isn't that he believes, on these occasions, that *he* is not appeared to redly; perhaps that is impossible. Let us concede, for purposes of argument, that necessarily, if, at **t**, **S** is appeared to redly, then **S** does not at **t** believe that he is not appeared to redly. Let us also concede that if a person is appeared to redly and pays attention to his phenomenal field (perhaps asking himself whether and how he is being appeared to) then if he is being appeared to redly, he believes that he is. These concessions are consistent with **S**'s being such that whenever he is appeared to redly, he believes that no one is ever thus appeared to; for **S**, we may add, does not, on these occasions, pay any attention to his phenomenal field. He does not ask himself whether he is being appeared to redly, or, indeed, whether he is being appeared to at all. He simply finds himself, under these conditions, believing that no one is ever appeared to redly. Let us add both that **S** always does his best to fulfill his epistemic duty, and that this bizarre belief of his is coherent with the relevant body of beliefs, whatever that may be. And now suppose **S** is appeared to redly, on a given occasion,

forming the belief that no one is ever appeared to redly. That belief satisfies the coherence requirement; furthermore, **S** is doing his epistemic duty in accepting it. Nevertheless the belief in question, clearly enough, has little by way of positive epistemic status for **S**.

A second example: Timothy is a promising young artist with an intense, indeed, inordinate admiration for Picasso. Waiting at a supermarket checkout, he idly picks up a copy of *The National Inquirer*, reading therein that Picasso, contrary to what most of us have always thought, was really an alien from outer space. Due to nonculpable gullibility and his overwhelming admiration for Picasso, Timothy then forms the belief that he too is really an alien from outer space. The rest of his beliefs readjust themselves so as to form a coherent pattern; we may add that Timothy is wholly and nonculpably unaware of the psychological mechanisms at work in the formation of this belief, thus flouting no epistemic duties in believing as he does. The belief in question, then, satisfies the coherence condition; furthermore, it has permissive justification for him. Nevertheless it has little by way of positive epistemic status for him. Even if, *per impossibile*, it turned out that Timothy really *is* an alien from outer space, he certainly does not know that he is. And the reason, in this case as in the last, is clear: Timothy holds the belief in question because of cognitive malfunction, because of noetic deficiency.

A final example, suggested by Timothy's plight. Suppose at **t** I am in Oxford and know that I am; I am just outside the gates of Balliol College, idly observing a small but noisy flock of gowned undergraduates on their way to Examination Schools. I believe that the walls of Balliol are behind me, that Broad Street is before me, that I am standing upon a sidewalk, and so on. I am paying no attention to my phenomenal field, and hold no beliefs about my experience or how I am being appeared to. My beliefs at **t** form a coherent system; each is coherent with the rest. Now imagine that I leave Oxford, taking the train to London. My experience then changes in the normal way; my visual, auditory, and kinesthetic experience at **t***, when I am on the train bound for London, is just what one would expect. Furthermore, I am devoted to my epistemic duty, and am doing my level best to satisfy my obligation to try to bring it about that I am in the right relation to the truth. Due to a sudden burst of radiation as we pass a nuclear dump, however, I undergo a cerebral accident resulting in cognitive malfunction: my beliefs are no longer responsive to my experience and revert to what they were at **t**, when

I was in Oxford. Due to this cognitive dysfunction, at **t*** I believe just what I did at **t**: that I am in Oxford, that a flock of noisy undergraduates is passing by, that I am standing on a sidewalk just outside the walls of Balliol, and the like. My beliefs at **t*** are coherent, for they are the very beliefs I coherently held at **t**, which by hypothesis formed a coherent system of beliefs. I am also, of course, doing my epistemic duty. Nevertheless, the belief that I am then in Oxford has, obviously enough, very little by way of positive epistemic status for me. I conclude that coherence is not sufficient for positive epistemic status—either by itself or in conjunction with aptness for epistemic duty fulfillment.

III. Reliabilism

I turn thirdly to the *reliabilist* account of positive epistemic status. Here we are faced with an embarrassment of riches: there are many reliabilist accounts[16], and considerations that apply to some do not apply to others. For the sake of definiteness, I shall select three reliabilist accounts for brief consideration, although I believe that what I say applies to some of the others as well.

A. Robert Nozick

According to Nozick, **S** knows that **p** if and only if four conditions are satisfied: (l) **p** is true, (2) **S** believes **p**, (3) if **p** were not true, **S** would not believe **p**, and (4) if **p** were true, **S** would believe **p**.[17] The fourth condition may be initially a bit puzzling. The idea is approximately this: if **p** were true and things a bit different from what in fact they are, **S** would (still) believe **p**. In terms of possible worlds: in the nearby worlds in which **p** is true, **S** believes **p**. If conditions (3) and (4) hold, then, says Nozick, **S**'s belief "tracks the truth".

This account of knowledge has several interesting features. It turns out, for example, that while I know

> (3) I am here in my study (on earth, back home in Indiana, nowhere near Alpha Centauri),

I do not know

> (4) I am not a brain in a vat on Alpha Centauri, serving as a subject in an experiment in which the experimenters give me the very experiences and beliefs I do in fact have just now.

I don't know (4) because it does not meet the third condition: it is false that if (4) were not true, I wouldn't believe it. (If (4) were not true I would be a brain in a vat with the very experiences and beliefs I do in fact have and thus would believe (4).) I know (3) but don't know (4), despite the fact that I can clearly see that (3) entails (4); on the view in question, knowledge is not closed (or closeable) under known entailment. Nozick says he takes skepticism seriously (pp. 197ff); on his account, the skeptic is quite correct in claiming that I do not know that I am not a brain in a vat being given the very beliefs and experiences I would have under more normal circumstances; nor do I know that I am not being similarly deceived by a malicious Cartesian demon. But where the global skeptic typically goes wrong, says Nozick, is in concluding that I don't know that I am in my study; this latter belief meets the four conditions for knowledge, even if the former does not, and despite the fact that the latter entails the former.

I think Nozick's concession to the skeptic is more apparent than real. On his account, even if I don't know that I am not a brain in a vat, I do know that I am not a brain in a vat and am home in my study. (If that conjunction were true, I would believe it; if it were false, it would be false by virtue of the falsehood of its second conjunct, in which case I wouldn't believe the conjunction.)[18] Wherever the skeptic holds that I do not know the denial of some momentous proposition (such as **I am being systematically and massively deceived by a malicious Cartesian demon**), Nozick agrees; but on his view while indeed I do not know the denial of the momentous proposition in question, I do know the conjunction of its denial with some commonplace proposition. On his view the correct reply to the skeptic who claims you don't know that you are not a brain in a vat on Alpha Centauri is "Well yes, perhaps I don't know that, but I *do* know that I am here in my study and am not a brain in a vat on Alpha Centauri." We could hardly blame the skeptic for feeling that with friends like that....

More to the present point, however, I think we can easily see that the four conditions Nozick lays down as necessary and sufficient for knowledge are not sufficient either for knowledge or for positive epistemic status. I shall give three examples. First, suppose I am a brain in a vat on Alpha Centauri. My captors are running a cognitive experiment; they give me most of the experiences and beliefs I would have if I were at home carrying on my normal life, so that most of

what I believe is absurdly false. Now a federally promulgated law for cognitive experiments, on Alpha Centauri, is that the subject must be given the true belief as to what the largest city on earth is; they therefore give me the belief that Mexico City is the largest city on earth. They also give me overwhelming evidence for the proposition that Mexico City is *not* the largest city on earth; I believe I have read many independent demographic surveys, many maps and atlases, all of which unite in declaring that *Cleveland* is the largest city on earth. I have no evidence at all for my belief that Mexico City is the largest. As things stand, then, the conditions for knowledge are met: where **M** is **Mexico city is the largest city on earth**, I believe **M**, **M** is true, in the nearby worlds where **M** is true I (still) believe **M**, and (because of that federally promulgated law) if **M** were not true I wouldn't have believed that it was. But surely I don't know **M**; in fact M has little if any positive epistemic status (beyond permissive justification) for me.

A second example: Our spaceship has landed on a small planet near Alpha Centauri. Tests indicate that the atmosphere, temperature and other conditions are propitious for human life; we confidently open the hatch and step out. We are immediately appeared to in the way that ordinarily (on earth) goes with perceiving a tiger at about 30 feet; naturally enough we form the belief that there is a tiger there. As it turns out, there is indeed a tiger (or an Alpha Centaurian tiger counterpart) there; but unbeknownst to us, Alpha Centaurian tigers emit a sort of radiation that makes them invisible to human beings. This planet's atmosphere, however, is suffused with a subtle sort of gas that causes earthlings to be appeared to in that characteristic tigerish fashion. Finally, tigers on this planet are attended by a certain parasite—one specific to tigers—that emits a kind of radiation in the absence of which, in the conditions that prevail on this planet, human beings are instantly rendered unconscious. Then Nozick's conditions for knowledge are met: our belief that there is a tiger there is true; in the nearby worlds where that proposition is true we believe it, and if it were not true we would not believe it (if it were not true we would have been rendered unconscious). But surely we don't know this proposition; indeed, it has little or no positive epistemic status for us. Nozick's conditions for knowledge are in fact met, but (so to speak) just by accident.

A final example: I suffer from a disease a symptom of which is the following condition. When a victim's retinas are irradiated with purple

light, he is never appeared to visually but is instead appeared to aurally; he seems to hear a certain tune. The disease has a further symptom: whenever a victim seems to hear that tune, he forms the belief that there is something purple in the neighborhood. On a given occasion my retinas are thus irradiated and I form the belief that there is a purple object nearby. This belief meets Nozick's conditions for knowledge; but once more, surely, it does not constitute knowledge and it has little by way of positive epistemic status for me. I conclude that Nozick's suggested necessary and sufficient conditions for knowledge are not in fact sufficient. (If you think you can always increase your knowledge by deducing consequences of what you know, then you will think these conditions aren't necessary either). Something more must be added; below (section IV) I shall suggest a promising area in which to look for the missing element.

B. Dretske

A second style of reliabilism sees reliability as a matter of *probability*; on this sort of account a person is said to know a true proposition if he believes it, and the right probability relations hold between that proposition and its significant others. As an example I shall consider Fred Dretske's interesting "informational theoretic" analysis of knowledge in *Knowledge and the Flow of Information*.[19]

This analysis goes as follows:

> (D_1) **K** knows that **s** is **F** = **K**'s belief that **s** is **F** is caused (or causally sustained) by the information that **s** is **F** (p. 86).

Two preliminary comments: Dretske is concerned, here, primarily or exclusively with *perceptual* knowledge; in particular the account is not designed to apply to such items of **K**'s *a priori* knowledge as that, say, 7 + 5 = 12. Secondly, the account is restricted to what Dretske calls "de re content" (p. 66); it is restricted, he says, to the kind of case where what **K** knows is a piece of information of or about **s**.

Now what sort of animal is this **information that s is F**? And what is it for a thing of that sort—presumably an abstract object or ensemble of abstract objects—to cause or causally sustain a belief? So far as I can see, Dretske says little by way of answer to the first question. What he does give are many examples of the sort **the information that s is F**. There is, for example, the information that Sam is happy, that the peanut is under shell number 3, that Susan

is jogging. One might say that these are *bits* of information, except for the fact that the term "bit" has been pre-empted for a *measure* of information. We are to think of information as being generated by or associated with states of affairs; and the amount of information generated by a given state of affairs depends upon the number and probability of the possibilities that state of affairs excludes. Suppose I throw a fair 64 sided die. The information that the die came up on a side numbered from 1 to 32 reduces the possibilities by a half and carries one bit of information; the knowledge that the die came up on a side numbered from 1 to 16 reduces the possibilities by another half and accordingly carries two bits; the information that side 3 came up reduces the original 64 possibilities to 1 and carries six bits of information. As you can guess from the example, if a piece of news reduces **n** (equally probable) possibilities to 1, then the amount of information displayed by that piece of news is log (to base 2) **n**. In the general case, where the possibilities involved need not be equi-probable (and where P(**A**) is the probability that a given possibility **A** is true) the amount of information generated by **A** is given by

(D_2) I(**A**) = log (1/P(**A**)), i.e., $-\log$ P(**A**).

Now there are deep problems and deep unclarities here. What are the relevant alternative possibilities for, for example, **Paul is jogging**? D_2 is applicable only where the possibilities involved are finite in cardinality; is that so for a noncontrived real life possibility such as **Paul is now jogging**? Do these alternative possibilities *have* appropriate probabilities? These are pressing questions for an account of this kind, and I don't believe there are even reasonably satisfactory answers to them. I shall not stop to argue that here, however, because the notion of the *amount* of information doesn't enter crucially into Dretske's account of knowledge. Nor need we know, for Dretskian purposes, just what information *is*; all we really need to know is what it is for a piece of information to *cause* or *causally sustain* a belief. Here the answer is disarmingly straightforward:

> Suppose a signal *r* carries the information that *s* is *F* and carried this information in virtue of having the property *F'*. That is, it is *r's being F'* (not, say, its being *G*) that is responsible for *r*'s carrying this specific piece of information. Not just any knock on the door tells the spy that the courier

> has arrived. The signal is three quick knocks followed by a pause and another three quick knocks... . It is the temporal pattern of knocks that constitutes the information-carrying feature (*F'*) of the signal. The same is obviously true in telegraphic communication.
>
> When, therefore, a signal carries the information that *s* is *F in virtue* of having property *F'*, when it is the signal's *being F'* that carries the information, then (and only then) will we say that the information that *s* is *F causes* whatever the signal's *being F'* causes (87).

So far, then, what we have is that a person **K** knows that s is **F** if and only if **K** believes that **s** is **F** and there is a signal **r** such that **r** has some property **F'** in virtue of which it carries the information that **s** is **F**; and **r**'s having **F'** causes **K** to believe that **s** is **F**. To simplify matters, suppose we drop the reference to the property **F'** of the signal by virtue of which it carries the information that **S** is **F**. What the analysis then boils down to is that **K** knows that **s** is **F** if and only if **K** believes that **s** is **F** and this belief is caused by a signal that carries the information that **s** is **F**. What we still need to know, then, is what it is for a signal to carry the information that **s** is **F**. This is given by

> (D_3) A signal *r* carries the information that *s* is *F* = The conditional probability of *s*'s being *F*, given *r* (and *k*) is 1 (but, given *k* alone, less than 1) (p. 65).

Now **k**, as Dretske explains, is the *background knowledge* of the receiver. D_3 must therefore be relativized to be accurate; a signal may carry the information that **s** is **F** relative to you but not to me. You already know that **s** is **F**; so the probability of **s**'s being **F** relative to your background knowledge is 1; no signal carries the information that **s** is **F** relative to you. I don't know that **s** is **F**; so any signal **r** which is such that the probability of s's being F on **r&k** (where **k** is my background information) equals 1 carries the information that **s** is **F** with respect to me. If you know that **s** is **F**, then no signal carries the information that **s** is **F** with respect to you; if I don't know that **s** is **F**, then any state of affairs carries that information with respect to me if its conjunction with what I do know entails that **s** is **F**. We can therefore rewrite D_3 as

(D_4) **r** carries the information that **s** is **F** relative to **K** iff P((**s is F**) | (**r&k**)) = 1 and P((**s is F**) | **k**) < 1.

And now we can say that

(D_5) **K** knows that **s** is **F** if and only if **K** believes that **s** is **F** and there is a state of affairs **r** such that (1) **r** causes **K** to believe that **s** is **F** and (2) P((**s is F**) | (**r&k**)) = 1 and P((**s is F**) | **k**) < 1.

We saw above that the problematic notion of the amount of information associated with a specific event or state of affairs can safely be ignored, since that notion plays no role in Dretske's final account of knowledge. But now we see that the same goes for any other specifically information theoretic concept; this analysis of knowledge, when spelled out, involves only the notions of probability, belief, and causation. Nothing specifically information theoretic seems to be involved.[20]

D_5, I think suffers from two sorts of deficiencies. In a way, the deepest problem, I believe, is that there is no currently available conception of probability that will serve Dretske's purposes. The probabilities in question are to be *objective* (p.55); so personalist and subjective accounts of probability will not be relevant. That leaves frequency, propensity and logical accounts of probability. On frequency accounts, however, there really isn't any such thing as the probability of such a specific singular proposition as **Paul is jogging** on given evidence; for on this conception a probability is a ratio between classes. What we face here is the dreaded problem of the single case: on frequency and propensity accounts, how can we go from genuine probability statements (**the probability of a Frisian's being a swimmer is .4**) to something like a probability for a singular proposition about some specific Frisian (who of course is a member of *many* reference classes in which the frequency of the attribute **being a swimmer** may differ wildly)? According to Reichenbach, such propositions have a "fictitious probability", which is to be arrived at by "direct inference" and can be thought of as a kind of estimate or *posit*. And the problem for Dretske's account here is that all of the ways suggested for making such a direct inference refer to what we know in such a way that the resulting 'probability' for a singular proposition is *relative to our knowledge*.[21] But then the 'probability' of a singular proposition might be different for me than it is for you,

since your knowledge might significantly differ from mine.

Although I can't argue this here, the same result will follow on propensity accounts of probability. And here is the problem for Dretske: since the probability of singular propositions, on both these conceptions, is relative in this way to bodies of knowledge, the same will be true, on Dretske's account of knowledge, for propositions ascribing knowledge to a person: these propositions will inherit that relativity. But then we have, on Dretske's account, the distressing result that relative to *me*, it may be true that Paul knows that *s* is *F*, but relative to *you*, false that he does. So neither frequency nor propensity accounts will be of use to Dretske.

This leaves only the *logical* account of probability; but this account, from Dretske's point of view, suffers not only from its own intrinsic implausibility but from a difficulty specific to Dretske's theory: the logical account conjoined with what Dretske says about probability entails that causal laws are necessary in the broadly logical sense.

It therefore seems that there is no conception of probability that will serve Dretske's needs. But (as he points out) his theory really requires only the limiting case where the probability of one proposition on another is 1; and perhaps (as he suggests) that notion can be replaced by the notion of "a particular kind of lawful dependency between signal and source .".[22] In any event, I don't have the space here to detail these objections; let me instead turn to an objection more germane to our present interests. This objection is just that even if we had a relevant conception of probability, Dretske's conditions for knowledge would not be sufficient. I shall give three examples to make the point. First, suppose **K** suffers from a serious abnormality—a brain lesion, let's say. This lesion wreaks havoc with **K**'s noetic structure, causing him to believe an array of propositions most of which are absurdly false. It also causes him to believe, however, that he is suffering from a brain lesion. **K** hasn't the slightest shred of evidence for this belief; and he thinks of his other unusual beliefs as resulting from no more than an engagingly original turn of mind. According to D_5, however, it follows that **K** *knows* that he suffers from a brain lesion. His having this lesion causes him to believe that he is thus affected; furthermore the probability of his suffering from a brain lesion on his background information alone is less than 1, but of course its probability on **k** and **K** is **suffering from a brain lesion** is 1. But surely **K** does *not* know that he is suffering from a brain lesion. He has no evidence of any kind—sensory, memory, in-

trospective, whatever—that he has such a lesion; his holding this belief is, from a cognitive point of view, no more than a lucky (or unlucky) accident. Indeed, we can add, if we wish, that **K** has powerful evidence for the conclusion that he is *not* thus suffering; he has just been examined by a trio of world famous experts from New York, who assure him that his brain is entirely normal. In this case, then, **K**'s belief that he has a brain lesion is not only such that he has no evidence for it; he has first rate evidence *against* it. In such a situation **K** clearly does not know that he has a brain lesion, despite the fact that this belief meets Dretske's conditions for knowledge.

Examples of this kind can be multiplied; so let's multiply a couple. You have wronged me; you have stolen my Frisian flag. By way of exacting revenge I sneak into your house at night and implant in your dog a source of extremely high frequency radiation. This radiation has no effect upon either you or your dog, except to cause you to form the belief that aliens from Alpha Centauri have invaded your house and replaced your dog with a non-terrestrial look-alike that emits ultraviolet radiation. You christen this creature (who is in fact your dog) 'Spot'. Your belief that Spot emits ultraviolet radiation then satisfies Dretske's conditions for knowledge: Spot's emitting ultraviolet radiation causes you to believe that he does; relative to what you know this is not probable, but relative to the conjunction of what you know with **Spot emits ultraviolet radiation**, its probability is, of course, 1. But surely you don't know that Spot emits such radiation. Indeed, as in the previous case we can add that you have powerful (though misleading) evidence **against** that proposition. You have had Spot examined by a highly competent group of physicists based at the Stanford linear accelerator; I have corrupted them, bribing them to tell you that Spot is entirely normal; but you are nevertheless unable to divest yourself of the belief in question. Surely you don't know.[23]

A final example: you and I each hold a ticket for a valuable lottery; the winner gets an all-expenses-paid week in Philadelphia. I approach the official drawer and get him to agree to fix the lottery: I am to coat my ticket with a substance **S** and he is to coat his hand with a substance **S*** in virtue of which my ticket will stick to his hand. After I leave, you appear and offer him twice as much; he accepts. He then coats his hand with a substance **S**** that causes *your* ticket to stick to his hand, thus causing you to win the lottery. It also causes me, by virtue of a cerebral abnormality on my part which is other-

wise undetectable, to believe that you will win. You and I witness the drawing; I suddenly and unaccountably find myself with the belief that you will win. On Dretske's account, I know that you will win, despite my knowledge that I have fixed the lottery. For (where T5 is your ticket) T5's being coated with **S**** causes me to believe that you will win; that you will win (let's suppose) has a probability of 1 on the conjunction of my background knowledge with **T5 has been coated with S**** but a vastly lower probability on my background knowledge alone. It is obvious, however, that under these conditions I don't know that you will win.

Clearly, there are as many examples of this sort as you please. One recipe for constructing them is just to consider some event **e** that causes **K** to believe that **e** occurs (or to believe some proposition entailed by **e**'s occurrence) where **e** causes **K** to form the belief in question by virtue of some pathological condition on **K**'s part—a brain lesion, let's say—and in such a way that it is a mere accident, from a cognitive point of view, that the belief is true. And what these examples show is that something further must be added to Dretske's account; the condition he suggests is not in fact sufficient. In section IV I shall make a suggestion as to what it is that must be added.

C. Goldman

Alvin Goldman suggests still another version of reliabilism, one that deserves that title *in excelsis*: "The justificational status of a belief", he initially says, "is a function of the reliability of the process or processes that cause it, where (as a first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false" (10)[24]. After some interesting preliminary skirmishes, he gives his official account in a sort of recursive form:

(a) If *S*'s belief in *p* 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.

(b) If *S*'s belief in *p* at *t* results ("immediately") from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing *S*'s belief in *p* at *t* are themselves justified, then *S*'s belief in *p* at *t* is justified (13, 20).

He then adds an appropriate closure clause.

For present purposes we need not concern ourselves with (b); suppose instead we turn our attention to the base clause (a). It isn't easy to see exactly what (a) comes to. The notion of **resulting from** is imprecise enough to require a good deal of guesswork, as is the notion of a process's being available to **S**. And while Goldman concedes that the second condition in the antecedent ("there is no reliable or conditionally reliable...") isn't quite accurate, he doesn't tell us how to set things right. Furthermore, this condition is subject to some of the familiar problems often bedeviling analyses involving counterfactuals. As an example of what requires this condition, Goldman cites a case in which I accept a lot of memory beliefs dating from my boyhood, but also have a lot of testimonial evidence from my parents—fabricated evidence, in point of fact—that my ostensible memories from that period are for the most part false. If nonetheless I persist in accepting the memory beliefs, then, says Goldman, these beliefs are not justified for me. They fall under that second condition in the antecedent of (a): there is available to me the process of using one's evidence, which is such that if I were to use it "in addition" to my memory, then I would not accept the beliefs in question. But now suppose that this process—using one's available evidence—is the only process which is at all likely to inhibit the relevant memory beliefs; and suppose further that if I were to use this process, then I would become suspicious of my parents, begin to question their veracity, launch an investigation into the case, conclude that my memories are reliable after all, and continue to accept them. Then the antecedent of (a) is fulfilled, for these beliefs; but they would not, presumably, be justified.

Still another problem here is the fact that there may be more than one reliable process available to **S**; perhaps there is a reliable process **P** such that if he were to use it, then he would not believe **p**, but another reliable process **P*** such that if he were to use both **P** and **P*** then he would believe **p**, and no reliable process **P**** such that if he were to use all of **P**, **P*** and **P**** he would not believe **p**; then perhaps Goldman would want to say that **S**'s belief would be justified after all, despite the fact that there is a reliable process available to him which is such that if he were to use it then he would not believe **p**.

So there are problems of several sorts with (a); I want to concentrate on only one. (a) proclaims that if **S**'s belief in **p** results from

reliable cognitive process (and meets that further condition), then S's belief is justified. But what are these "cognitive processes"? Here is Goldman's way of characterizing them:

> Let us mean by a 'process' a *functional operation* or procedure, ie., something that generates a *mapping* from certain states—'inputs'—into other states—'outputs'. The outputs in the present case are states of believing this or that proposition at a given moment. On this interpretation, a process is a *type* as opposed to a *token*. This is fully appropriate, since it is only types that have such statistical properties as producing truth 80% of the time; and it is precisely such statistical properties that determine the reliability of a process. Of course we also want to speak of a process as *causing* a belief, and it looks as if types are incapable of being causes. But when we say that a belief is caused by a given process, understood as a functional procedure, we may interpret this to mean that it is caused by the particular *inputs* to the process (and by the intervening events 'through which' the functional procedure carries the inputs into the output) on the occasion in question (ll).

How, exactly, shall we understand this? A belief forming process, first of all, is a thing that "generates a mapping from certain states—'inputs'—into other states—'outputs'." Now a barometer or a thermometer generates such a mapping: a thermometer, for example, generates (among others) a mapping from the ambient temperature to the reading it displays. Goldman's processes, however, are not concrete instruments or mechanisms such as barometers, thermometers, anometers, sphygmomanometers and the like; he says they are *types*, and he says that these types themselves have inputs and outputs. That makes it sound very much as if the processes in question are *functions*—but if they were functions they would not merely *generate* mappings (as he says they do), they would *be* mappings. So what sort of types are these? What are their tokens like? You might think that the tokens of the relevant types would be such things as Paul's memory, or Paul's visual apparatus. These things would not be functions, or abstract objects of any kind, but specific and concrete belief producing mechanisms or faculties. Like functions, such mechanisms or faculties have an input and output; but

they are analogous to barometers and altimeters rather than to functions from temperature or altitude to numbers. The problem with this plausible answer, however, is that according to Goldman it is only the relevant **types**, not their tokens, that are reliable or unreliable. I'm not sure why Goldman said this and I wish he hadn't; I should think a concrete instrument or mechanism—a barometer, for example—is quite properly said to be reliable or unreliable.

So how shall we think of these processes? What are their tokens like? Do their tokens also have inputs and outputs? What sorts of things are the inputs and outputs of these processes, or the inputs and outputs of their tokens (if their tokens have inputs and outputs)? It is clear, I think, that the *outputs* of these processes are beliefs; the inputs, apparently, are entities capable of *causing* the outputs. But then presumably the beliefs in question will not be *propositions* (propositions aren't caused by anything) but *events*: such events as **Paul's believing that all men are mortal**, for example. Presumably, then, the inputs will also be events, rather than such abstract objects as states of affairs or properties or types—such events, perhaps, as **Paul's mother telling him that all men are mortal**. Goldman gives some examples of the sorts of processes he has in mind:

> One example is reasoning processes, where the inputs include antecedent beliefs and entertained hypotheses. . . . A third example is a memory process, which takes as input beliefs or experiences at an earlier time and generates as output beliefs at a later time. For example, a memory process might take as input a belief *at* t_1 that Lincoln was born in 1809 and generate as output the belief *at* t_n that Lincoln was born in 1809 (11-12).

Consider the memory process (the type) mentioned. The inputs for this process would be events consisting in a person **S**'s believing at a time **t** that Lincoln was born in 1809: **Paul's believing at t_1 that Lincoln was born in 1809, Sam's believing at t_2 that Lincoln was born in 1809**, etc.; the outputs for a given input will be an event consisting of **S**'s believing that same proposition at a later time: **Paul's believing at t_3 that Lincoln was born in 1809, Sam's believing at t_4 that Lincoln was born in 1809**, and the like. For any input or argument **x** of the process, furthermore, the value or output **x*** will be caused by **x**. Tokens of this type, furthermore, would be specific dateable sequences of concrete events—events taking place

in Paul's brain, perhaps, but in any case events taking place somewhere in his cognitive apparatus. The types would be types of such tokens. Both tokens and types will have inputs and outputs; an input (for a token as well as a type) will be some event or state of the cognizing subject, and an output will be a specific concrete event of the sort consisting in Paul's believing at some time t that Lincoln was born in 1809. And a given input (whether token or type) will be causally involved in a relevant way with its corresponding output.

Plenty of problems about the ontology of these processes remain: for example, a given token of a process will be a sequence whose members are concrete events, and presumably concrete events in the cognitive apparatus of a single person. But what about the types? Will they too then be such items as **the type consisting in e_1, followed by e_2, ..., followed by e_n** (where 'e_1', 'e_2', etc. are *names* of specific concrete events)? Or would the types be such that their tokens were sequences of *types* rather than sequences of concrete events: $e^*_1, e^*_2, ..., e^*_n$, where each e^*_i is a relevant type of e_i? If the latter, would these types involve a reference to a particular person? Suppose we temporarily ignore these questions, saving them for a more propitious occasion. The main problem, as I see it, still remains. Note first that any particular token—any relevant sequence of concrete events—will be a token of many different types. Consider a specific visual process in Paul, where the input consists in retinal stimulation, let's say, and the output consists, for some particular scene **s** on his television, in his believing that he sees **s**. The process in question will presumably involve a large number of events; it will no doubt include an event consisting in Paul's being appeared to in a characteristic way. Now this sequence of events will be a token of many different types—**the cognitive process, the visual process, the cognitive process occurring on a Thursday, the visual process occurring in a middle aged man, the visual process occurring in a middle aged man under such and such lighting conditions, the visual process occurring in a middle aged man when his retinas are being stimulated by light of such and such a character**, and many more.

It is these types that are to be evaluated for reliability (since, as we recall, the degree of justification enjoyed by the belief in question is a function of the reliability of the process (type) causing it); but obviously the types may differ wildly among themselves with respect to reliability. Which is the relevant type? Which type is the

one such that its reliability determines the justification Paul has for the belief in question? This is the *problem of generality*, noted by Goldman (12) and developed by Richard Feldman in "Reliability and Justification".[25] Now obviously we can't take the relevant type to be, say, **the cognitive process**, or **vision**; for the outputs of such processes will have many different degrees of justification. If the reliability of a given belief— Paul's belief that he is watching television, for example—is to be determined by the reliability of the relevant type of which it is an output, then the outputs of that type must be indistinguishable with respect to justification: they must have the same degree of justification, or, if we eschew degrees in favor of a comparative concept of justification, none must be more justified than any other. Furthermore, the relevant type must display the degree of reliability correlated with the (no doubt modest) degree of justification Paul's belief does in fact have. Still further, the type in question clearly couldn't be such a type as **cognitive process issuing in a true belief**; there can't be a specific degree of reliability such that it is a necessary truth that the type in question displays that degree of reliability.

These things mean that the relevant types must display a very considerable degree of specificity. Consider, for example, (to take a case like one of Goldman's) an occasion on which Paul sees a mountain goat. Here we have a sequence of concrete events, a sequence that takes something (retinal stimulation, e.g.) as input and as output yields the event of Paul's forming the belief: **that's a mountain goat**. Now what is the relevant type, the type whose reliability determines the degree of justification Paul has for this belief? Not, of course, **vision**, or **vision in the mountains**, or **vision on Mt. Shuksan** or **vision during daylight**, but something much more specific: **seeing a mountain goat at 350 yards under such and such light and atmospheric conditions**, perhaps, or **retinal stimulation of such and such a character under such and such lighting and atmospheric conditions**[26], (where it might be very hard to fill in the such and such's). Even this won't be specific enough, however; for people differ with respect to their familiarity with mountain goats, liability to buck fever, visual acuity at 350 yards, and the like. These process types, therefore, are of differing degrees of reliability for different people and hence yield differing degrees of justification for different people; we must add further parameters to the type in question.

Indeed, we might plausibly think that for any belief **B**, type **T** and specification **T*** of **T** (where, for example, the type **night vision** is a specification of the type **vision**) if **p** is in the output of both **T** and **T*** and **T** and **T*** differ with respect to reliability, then **T** will not be the type the reliability of which determines the justification of **B**. A type **T**, we might think, will be a relevant type for a belief **B** of Paul's—relevant in that its reliability determines the justification of **B**—only if it is **maximally specific with respect to reliability** ('max' for short), i.e., such that there is no specification of that type with a different degree of reliability. I say this is a plausible thought; but so far as I can see it does not follow from what Goldman says. (Of course it would be hard to see grounds for its rejection; on Goldman's view, what determines the degree of justification of a belief is the reliability of the relevant type generating it; and it is hard to see grounds for supposing that there might be a pair of types **T** and **T***—both, let us say, pertaining to psychological conditions alone—such that **T*** is a subtype, a specification of **T**, while it is **T** rather than **T*** the reliability of which determines the degree of justification of beliefs in the output of **T***.)

In any event, although it would be difficult to give an example of the types in question (i.e., the types the reliability of which determine the justification of the beliefs in their outputs), it is easy to see that they will have to be types of very great specificity. And we can easily see further, I think, that the degree of justification of a belief issued by such a type and the degree of reliability enjoyed by that type will not nearly always be related in the way required by Goldman's theory. The cases that caused trouble for Dretske will also cause trouble, and similar trouble, for Goldman. Consider, for example, the person whose belief that he has a brain tumor is caused by his brain tumor. There is a rare but specific sort of brain tumor, we may suppose, such that associated with it are a number of cognitive processes of the relevant degree of specificity, most of which cause its victim to hold absurd beliefs. One of the processes associated with the tumor, however, causes the victim to believe that he has a brain tumor. Suppose, then, that **S** suffers from this sort of tumor and accordingly believes that he suffers from a brain tumor. Add that he has no evidence at all for this belief: no symptoms of which he is aware, no testimony on the part of doctors or other expert witnesses, nothing. Then the relevant type, while it may be hard to specify in detail, will certainly be highly reliable; but surely it is

not the case that this belief—the belief that he has a brain tumor—has much by way of positive epistemic status for **S**. Indeed, as in the Dretske case, we can add, if we like, that **S** has a great deal of evidence *against* the proposition that he has a brain tumor; he too has just been examined by that team of brilliant specialists from New York, who have given him a clean bill of health. This addition does not run afoul of the second clause of the antecedent of (a) above because **S** is such that if he were to use the "process" of consulting his available evidence (the only other reliable process available to him), he would become very much interested in the whole question of tumors, study the matter in considerable detail, make some new discoveries, and finally wind up concluding that he did indeed have a tumor. There is therefore no reliable process available to **S** which is such that if he had used it, then he would not have formed the belief that he has a tumor. The antecedent of (a) is then satisfied, for **S**[27]; but surely the belief in question has little positive epistemic status for him. Perhaps it has permissive justification; perhaps **S** is within his epistemic rights in holding the belief; indeed, perhaps it is not within **S**'s power *not* to accept this belief. But the belief in question has little or no positive epistemic status of any other sort for him; in particular, it is clear that it does not constitute *knowledge* for him, even though it is true and he is permissively justified in accepting it.

Obviously we can construct similar examples from the other cases that caused trouble for Dretske: the case where I implant a source of high energy radiation in your dog, for example, and the case of the doubly rigged Philadelphia lottery. As a result of cognitive malfunction, the degree of positive epistemic status enjoyed by a belief **B** does not, in these cases, match the reliability of the relevant process type that produces it. And obviously the general recipe for constructing such cases is to come up with an appropriately pathological process type of the right degree of generality which is in fact reliable, but (due to the pathology involved) does not confer much by way of positive epistemic status on the beliefs in its output. (Such types, we might say, are from a cognitive point of view *accidentally reliable*.[28]) Reliability, then, is not sufficient for positive epistemic status. A further condition must be added—a condition to which it is now time to turn.

IV. On Working Properly

A. The Basic Idea

In the above discussions there is a sort of recurring theme. We saw repeatedly that various proposed analyses of justification come to grief when we reflect on the variety of ways in which our noetic faculties can fail to function properly. Chisholm's dutiful epistemic agent who, whenever he is appeared to redly, always believes that nothing is appearing redly to him, Lehrer's coherent epistemic agent who believes that he is an alien from outer space, Dretske's and Goldman's epistemic agent whose belief that Spot emits ultraviolet radiation has been caused by the fact that Spot does indeed emit such radiation—all are such that their beliefs lack positive epistemic status for them. In each case the reason, I suggest, is *cognitive malfunction*, failure of the relevant cognitive faculties to function properly. Chisholm's agent believes as he does because of cognitive dysfunction due to a brain lesion, or the machinations of an Alpha Centaurian scientist, or perhaps the mischievous schemes of a Cartesian evil demon; and something similar can be said for the others. I therefore suggest that a necessary condition of positive epistemic status is that one's cognitive equipment, one's belief forming and belief sustaining apparatus, be free of such cognitive malfunction. A belief has positive epistemic status for me only if my cognitive apparatus is functioning properly, working the way it ought to work in producing and sustaining it.

It is important to see is that this condition—that of one's cognitive equipment functioning properly—is not the same thing as one's cognitive equipment functioning *normally* in the statistical sense. If I give way to wishful thinking, forming the belief that I will soon be awarded a Nobel Peace Prize, then my cognitive faculties are not working properly—even though wishful thinking may be widespread among (and in that sense normal for) human beings. It may be (and in fact is) the case that it is not at all unusual or abnormal for a person to form a belief out of jealousy, lust, contrariness, desire for fame, or wishful thinking; nevertheless when I form a belief in this way my cognitive equipment is not functioning properly; it is not functioning the way it ought to.

Suppose we initially and provisionally agree that a necessary condition of a belief's having positive epistemic status for me is that the relevant portion of my noetic equipment involved in its formation

and sustenance be functioning properly. Still, your faculties' being in proper working order cannot be the whole story. You take a space trip to a planet revolving about a distant star—Alpha Centauri, say. Conditions there are much like they are on earth; indeed some of the inhabitants of the planet are (physiologically speaking) surprisingly similar to human beings. Conditions there are propitious for human life; still there are subtle epistemic differences. Cats (or their Alpha Centaurian counterparts) are invisible to human beings; but they emit a sort of radiation unknown on earth, a radiation which works directly on the appropriate portion of a human brain, causing its owner to form the belief that a dog is barking nearby. An Alpha Centaurian cat slinks by; you form the belief that a dog is barking nearby. There is nothing the matter with your noetic faculties, but the belief in question has very little positive epistemic status for you. The problem is not merely that the belief is false; even if a dog is barking nearby (but in a soundproof room, say, so that it is inaudible), you certainly don't know that it is. The problem is that your cognitive faculties and the environment in which you find yourself are not properly attuned. The problem is not with your cognitive faculties: they are in good working order; the problem is with the environment. In much the same way, your automobile might be in perfect working order, despite the fact that it will not run well at the top of Pike's Peak, let alone under water. So we must add another component to positive epistemic status; your faculties must be in good working order, and the environment must be appropriate for your particular repertoire of epistemic powers. (Cats are not invisible to the human like creatures native to the planet in question.)

We might think that a belief's having positive epistemic status just *is* its being produced by epistemic faculties that are functioning properly (in producing and sustaining that belief) in an epistemically appropriate environment; but this cannot be the whole story. For couldn't my cognitive faculties could be working properly (in an appropriate environment) in producing and sustaining a certain belief in me, while nonetheless that belief has little by way of positive epistemic status for me? Further, a pair of my beliefs could be **productively equivalent** (produced by faculties functioning properly to the same degree and in environments of equal appropriateness) but nonetheless such that one of them has vastly more by way of positive epistemic status than the other. *Modus ponens* (more exactly, its corresponding conditional) obviously has more by way of positive

epistemic status for me than for the memory belief, now rather dim and indistinct, that forty years ago I owned a second hand 16 gauge shotgun and a red bicycle with balloon tires. So positive epistemic status can't be simply a matter of a belief's being produced by faculties working properly in an appropriate environment.

What more is required? Here, fortunately enough, there is an easy response. Not only does the first belief, the belief in the corresponding conditional of *modus ponens*, have more by way of positive epistemic status for me than the second, it is also one I accept much more firmly; I have a much stronger tendency or inclination to accept that proposition than to accept the other. Obviously another element of positive epistemic status is the degree to which I do or am inclined to accept the belief in question; I can't be said to *know* **p**, for example, unless I believe it very firmly indeed. If my faculties are working properly, the more strongly I believe (or am impelled or inclined to believe) **p** the more positive epistemic status **p** has for me. When our cognitive establishment is working properly, the strength of the impulse towards believing a given proposition (we may conjecture) will be proportional to the degree it has of positive epistemic status—or if the relationship isn't one of straightforward proportionality, the appropriate functional relationship will hold between positive epistemic status and this impulse. So when my faculties are functioning properly, a belief has positive epistemic status to the degree that I find myself inclined to accept it; and this (again, if my faculties are functioning properly) will be the degree to which I *do* accept it.

Initially and to a first approximation, therefore, we may put it like this (thus importing what is at this stage undoubtedly a spurious precision): a belief **B** has positive epistemic status for **S** if and only if that belief is produced in **S** by his epistemic faculties working properly; and **B** has more positive epistemic status than **B*** for **S** iff **B** has positive epistemic status for **S** and either **B*** does not or else **S** is more strongly inclined to believe **B** than **B***. (If we think degree of belief and degree of positive epistemic status are real valued functions, we can put the matter thus: where a person **S** accepts a proposition **P**, **S** has positive epistemic status to degree **r** for **P** if and only if his faculties are functioning properly in producing this belief and **S** accepts **P** to degree **r***, where **r*** is the value for **r** under the appropriate functional transformation.)

Now I mean to suggest that this or something like it is in fact how we think of positive epistemic status. But of course if that is so, then

something else must also be so: we must also think, as we do think, that (for the most part) when our faculties are functioning properly, the beliefs they produce—in particular the more confident beliefs they produce—are true, or at any rate *close* to the truth.

Of course so far this is merely programmatic, just a picture. Much must be said by way of clarification, articulation, qualification. Let me begin with a couple of obvious qualifications. First, I say that a belief has positive epistemic status for me only if my epistemic faculties are working properly in producing and sustaining it. But of course it isn't true that *all* of my cognitive faculties have to be functioning properly in order for a belief to have warrant for me. Suppose my memory plays me tricks; obviously that doesn't mean that I can't have warrant for such introspective propositions as that I am appeared to redly. What must be working properly are the faculties involved in the production of the particular belief in question. And even they need not be working properly over the entire range of their operation. Suppose I cannot properly hear high notes: I may still learn much by way of the hearing ability I do have. Furthermore, a faculty that doesn't function properly *without outside aid* can nonetheless furnish warrant; I can have warrant for visual propositions even if I need glasses and can see next to nothing without them. Still further, even if my corrected vision is very poor, I can still have warrant for visual propositions; even if I can't distinguish color at all, I can still have warrant for the proposition that I perceive something round. Still further, even if I can't perceive colors at all, I can still have visual warrant for the proposition that something is red; even if for me nothing appears redly (everything is merely black and white) I might still be able to see that something is red, in the way in which one can see, on a black and white television, which boxer is wearing the red trunks. Similar comments must be made, of course, about the environmental condition. There are problems here, but nothing that looks initially insurmountable.

Further, proper functioning, of course, *comes in degrees*; or if it does not, then approximation to proper functioning does. Clearly the faculties relevant with respect to a given belief need not be functioning *perfectly* for me to have warrant for my belief; how well, then, must they be functioning? And precisely how similar to the environment for which my faculties were designed, must my present environment be if I am to have warrant? Part of the answer here, of course, is that there is no answer; the ideas of warrant and knowledge are to some degree vague; there is therefore no precise

answer to the question in question.

Still further, suppose I know that the environment is misleading; and suppose I know in just which ways it is misleading. (I'm on a planet where things that look square are really round.) Then, clearly enough, the fact that my environment is misleading need not deprive my beliefs of warrant. And of course the same must be said for the requirement that my faculties be in good working order. If I know, for example, that (as in Castaneda's fantasy) I suffer from a quirk of memory such that whenever I read a history book, I always misremember the dates, somehow adding ten years to the date as stated, I can still have warrant for beliefs—even beliefs about dates—formed by reading history books; I can compensate for my erroneous tendency. What counts here, of course, are uncorrected and uncompensated malfunctionings. Clearly there is here the need for a great deal of chisholming; let me postpone it, however, in order to turn to some other matters.

B. The Design Plan

A crucially important notion here is that of specifications, or blueprint, or *design plan*. We take it that when the organs (or organic systems) of a human being (or other organism) function properly, they function *in a particular way*. Such organs have a *function* or *purpose*; the purpose of the heart, for example, is to pump blood. Furthermore, such an organ, of course, functions in such a way as to fulfill its purpose; but it also functions to fulfill that purpose in just one of an indefinitely large number of possible ways. Here a comparison with artifacts is useful. A house is designed to produce shelter—but not in just any old way. There will be plans specifying the length and pitch of the rafters, what kind of shingles are to be applied, the kind and quantity of insulation to be used, and the like. Something similar holds in the case of us and our faculties; we seem to be constructed in accordance with a specific set of plans. Better (since this analogy is insufficiently dynamic) we seem to have been constructed in accordance with a set of specifications, in the way in which there are specifications for, for example, the 1983 GMC van. According to these specifications (here I am just guessing), after a cold start the engine runs at 1500 RPM until the engine temperature reaches 140 degrees F.; it then throttles back to 750 RPM.

In the same sort of way, there is something like a set of specifications for a well-formed, properly functioning human being—an ex-

traordinarily complicated and highly articulated set of specifications, as any first year medical student could tell you. Suppose we call these specifications a 'design plan', leaving open the question whether human beings and other creatures have in fact been designed. Then of course the design plan will include specifications for our cognitive faculties (as well as for the rest of our powers and faculties). They too can work well or badly; they can misfunction or function properly. They too work in a certain way when they are functioning properly—and work in a certain way to accomplish their purpose. The purpose of the heart is to pump blood; that of our cognitive faculties is to supply us with reliable information— information about our environment, about the past, about the thoughts and feeling of others, and so on. But not just any old way of accomplishing this purpose in the case of a specific cognitive process is in accordance with our design plan. It is for this reason that it is possible for a belief to be produced by a cognitive process or belief producing mechanism that is *accidentally* reliable, as in the case of the processes cited as counterexamples to Goldman's version of reliabilism.[29] Although the belief producing processes in question are in fact reliable, their outputs have little by way of positive epistemic status; and the reason is that these processes are pathologically out of accord with the design plan for human beings.

Our design plan, of course, is such that our faculties are highly responsive to circumstances. Upon considering an instance of *modus ponens*, I find myself believing its corresponding conditional; upon being appeared to in the familiar way, I find myself holding the belief that there is a large tree before me; upon being asked what I had for breakfast, I reflect for a moment and then find myself with the belief that what I had was eggs on toast. In these and other cases I do not *decide* what to believe; I don't total up the evidence (I'm being appeared to redly; on most occasions when thus appeared to I am in the presence of something red; so most probably in this case I am) and make a decision as to what seems best supported; I simply find myself with the appropriate belief. Of course in some cases I may go through such a weighing of the evidence; I may be trying to decide, for example, whether the alleged evidence in favor of the theory that human life evolved by means of the mechanisms of random genetic mutation and natural selection from unicellular life (which itself arose by substantial similar mechanical processes from nonliving material) is in fact compelling; but in the typical case of belief formation nothing like this is involved.

According to our design plan, obviously enough, *experience* plays a crucial role in belief formation—both sensuous experience, such as **being appeared to greenly**, and the sort of experience involved in feeling impelled or disposed to accept a given belief. *A priori* beliefs, for example, are not, as this denomination mistakenly suggests, formed prior to or in the absence of experience. Thinking of the corresponding conditional of **modus ponens** *feels* different from thinking of, say, the corresponding conditional of **affirming the consequent**; and this difference in experience is crucially connected with our accepting the one and rejecting the other. Of course experience plays a different role here from the role it plays in the formation of perceptual beliefs; it plays a still different role in the formation of memory beliefs, moral beliefs, beliefs about the mental lives of other persons, beliefs we form on the basis of inductive evidence, and the like. Further, our design plan is such that under certain conditions we form one belief *on the evidential basis* of others; and of course if our faculties are functioning properly, we don't form just *any* belief on the evidential basis of just any other. I may form the belief that Sam was at the party on the evidential basis of other beliefs—perhaps I learn from you that Sam wasn't at the bar and from his wife that he was either at the bar or at the party. But if my faculties are functioning properly, I won't form the belief that Feike is a Catholic on the evidential basis of the propositions that nine out of ten Frisians are Protestants and Feike is a Frisian. And here too experience plays an important role. The belief about Sam *feels like* the right one; that belief about Feike (in those circumstances) feels strange, rejectable, inappropriate, not to be credited. Still further, the design plan dictates the appropriate *degree* or firmness of a given belief in given circumstances. You read in a relatively unreliable newspaper an account of a 53 car accident on a Los Angeles freeway; perhaps you then form the belief that there was a 53 car accident on the freeway. But if you hold that belief as firmly as, e.g., that $2+1=3$, then your faculties are not functioning as they ought to and the belief has little positive epistemic status for you. Again, experience obviously plays an important role. What we need is a full and appropriately subtle and sensitive description of the role of experience in the formation and maintenance of all these various types of beliefs; that project will have to await another occasion, as one says when one can't in fact deliver the goods.

Positive epistemic status, I said, is intimately related to proper function—so intimately related that a belief has positive epistemic

status for me only if my cognitive faculties are functioning properly in forming and maintaining it. As I said above, a presupposition of our conception of positive epistemic status, clearly enough, is that beliefs formed by our faculties functioning properly in an appropriate environment will be for the most part close to the truth. We might put this by saying that a presupposition of our conception is that the purpose of our epistemic faculties is the production of beliefs that are mostly true or mostly nearly true. Still, there are cases in which our faculties are functioning perfectly properly, but where their working in that way does not seem to lead to truth—indeed, it may lead away from it. Someone may remember a painful experience as less painful than it was, as is sometimes said to be the case with childbirth. You may continue to believe in your friend's honesty long after evidence and objective judgment would have dictated a reluctant change of mind. I may believe that I will recover from a dread disease much more strongly than the statistics justify. In these cases, the relevant faculties may be functioning properly, functioning just as they ought to, but nevertheless not in a way that leads to truth, to the formation of true beliefs. But then how can I say that a belief has positive epistemic status if it is produced by one's faculties functioning properly?

The answer here is simplicity itself. Different parts or aspects of our cognitive apparatus could have different purposes; different parts or aspects of our design plan, could be aimed at different ends or goals. Not all aspects of the design of our cognitive faculties need be aimed at the production of true belief; some might be such as to conduce to survival, or relief from suffering, or the possibility of loyalty, or inclination to have more children, and the like. What confers positive epistemic status is one's cognitive faculties working properly, or working according to the design plan *insofar as that segment of the design plan is aimed at producing true beliefs.* But someone whose holding a certain belief is a result of an aspect of our cognitive design that is aimed not at truth but at something else won't properly be said to know the proposition in question, even if it turns out to be true (unless, of course, the same design would conduce both to truth and to the other state of affairs aimed at.)

Finally, our design plan obviously dictates change over time; our faculties and organs change and mature. Newborn babies are not able to walk; kittens are. Kittens are born blind; human beings are not, although there is maturation of human cognitive faculties just as of those of other organisms. Still further, the design plan of an

organism may itself change over time, so that, conceivably, what is proper function at one time may not be at another. And of course the present view involves no specific or species chauvinism; it isn't necessary that a member of another species—an angel, for example, or an Alpha Centaurian—will have positive epistemic status for her beliefs only if her cognitive faculties function in accordance with *our* design plan.

C. Gettier Problems

I began this paper with a reference to Edmund Gettier and the salutary havoc his three page paper has introduced into contemporary epistemology. Gettier pointed out, of course, that belief, truth and justification are not sufficient for knowledge. Naturally enough, there have been many attempts to provide a "fourth condition", many attempts to add an epicycle or two to circumvent Gettier. Sadly, however, in most cases the quick response has been another Gettier problem that circumvents the circumvention. I don't mean at all to denigrate this often illuminating literature; but my aim here is not to enter the lists and try to produce a Gettier-proof analysis of knowledge. My aim instead is to see how the Gettier problem looks from the vantage point of the present conception of positive epistemic status. Gettier problems come in several forms. There is the original **Smith owns a Ford or Brown is in Barcelona** version: Smith comes into your office, bragging about his new Ford, shows you the bill of sale and title, takes you for a ride in it, and in general supplies you with a great deal of evidence for the proposition that he owns a Ford. Naturally enough you believe the proposition **Smith owns a Ford**; acting on the maxim that it is good to believe as many truths as possible, you infer from that proposition its disjunction with **Brown is in Barcelona**, where Brown is an acquaintance of yours about whose whereabouts you have no information. As it turns out, Smith is lying (he does not own a Ford) but Brown, by sheer coincidence, is indeed in Barcelona. So your belief **Smith owns a Ford or Brown is in Barcelona** is indeed both true and justified; but surely you can't properly be said to *know* it. A similar case, due to Lehrer: you see (at about 50 yards) what you think is a sheep in the field and (acting again on the above principle) infer that the field contains at least one sheep. As luck would have it, what you see is not a sheep (but a wolf in sheep's clothing); in a part of the field you can't see, however, there is indeed a sheep. Your belief that there is a sheep

in the field is true and justified, but hardly a case of knowledge.

In these cases you infer the justified true belief from a justified false belief (that Smith owns a Ford, that *that* is a sheep); your justification, we might say, goes through a false belief. But of course this is not the key to Gettier problems. Consider the following case, due originally to Carl Ginet. You are driving through southern Wisconsin, near Waupun. In an effort to make themselves look more prosperous than they really are, the inhabitants have erected a large number of barn facades— three for each real barn. From the road, these facades are indistinguishable from real barns. You are unaware of this playful attempt at deception; looking at what is in fact a real barn you form the belief **that's a fine barn!** Again, the belief is true and it seems that you are justified in holding it; but surely it does not constitute knowledge. To continue the bucolic motif, one final case. The Park Service has just cleaned up a popular bridle trail in Yellowstone, in anticipation of a visit from a Department of the Interior bigwig. A wag with a perverse sense of humor comes along and scatters two bushels of horse manure on the trail. The official from the Department of the Interior arrives, goes for a walk on the trail, and forms the belief that horses have been on the trail recently. Once more, his belief is true and justified, but does not constitute knowledge.

But why not, precisely? What is going on in these cases? First, it seems that in each of these cases it is merely *by accident* that the justified true belief in question is true. It just happens that Brown is in Barcelona, that there is a sheep in another part of the field, that what you are looking at is a barn facade rather than a barn. In each of these cases, the belief in question could just as well have been false. (As a matter of fact, in some of these cases that's not putting it strongly enough; the belief could *better* have been false: there are so many other places Brown could have been, and wags don't often or ordinarily take the trouble to make the Park Service look bad.) But what is the force, here of saying that the beliefs are true by *accident*?

Here is a possibility. In each of these cases there is a sort of glitch in the cognitive situation, a minor infelicity due, we might say, to cognitive environmental pollution. We saw above that a necessary condition of my beliefs' having positive epistemic status is that the environment in which they are formed be appropriate for one with my repertoire of cognitive powers. If there is substantial lack of match

between the cognitive environment and the sort of environment for which my powers are appropriate, then even if my belief happens to be true, it has little by way of positive epistemic status for me; the cognitive environment is deeply misleading, so that if I acquire a true belief, it is just by accident. Now in Gettier situations we have a sort of mild version of what goes on in those cases where there is wholesale lack of match between cognitive environment and cognitive faculties. In the Gettier cases there is no wholesale lack of match; there is, however, a bit of retail lack of match. Although the cognitive environment is not deeply misleading, it is nonetheless at least mildly misleading. Our design plan leads us to believe what we are told by others; there is what Thomas Reid calls "Credulity", a belief forming process whereby for the most part we believe what our fellows tell us. Of course Credulity is modified by experience; we learn to believe some people under some circumstances and disbelieve others under others. (We learn not to form beliefs about a marital quarrel until we have heard from both parties.) Still, Credulity is part of our design plan. But of course it doesn't work well when our fellows lie to us or deceive us in some other manner, as in the case of Smith who lies about the Ford, or the Wisconsinites who set out to deceive the city slicker tourists, or the wag aiming to hoodwink the Interior Department official.

We might generalize the idea of a design plan: there is a design plan not only for our cognitive faculties, but for the entire cognitive situation. Take the metaphor in this notion of design more seriously for the moment; then the designer of our cognitive powers will have designed those powers to produce mostly true beliefs in the sorts of situations their owners ordinarily encounter. The designer will be aiming at a kind of match between cognitive powers and cognitive environment; there will be, we might say, a sort of design plan not just for cognitive faculties, but for cognitive-faculties-cum-cognitive-environment. In Gettier situations, however, there are relatively minor departures from the design plan for the cognitive situation in question; the cognitive environment then turns out to be misleading for someone with our cognitive powers. And the force of saying that in these cases the beliefs just *happen* to be true, are true *by accident*, is the same as in the case of the counter-examples to reliabilism: the belief's being true is not a result of things working in accordance with the design plan. In the first cases, the problem was with the cognizer's faculties; due to disease or demon they were not function-

ing in accordance with the design plan. In typical Gettier situations, on the other hand, there is deviation from the design plan for the total cognitive situation, but it is due to the cognitive environment rather than to the cognizer's faculties. But this is a nonessential feature of Gettier situations; no doubt we could think of Gettier situations in which the glitch was internal to the cognizer rather than due to the environment. What is essential to Gettier situations is the production of a true belief despite a relatively minor failure of the cognitive situation to match its design.

D. Theism and Proper Function

But aren't such ideas as that of working properly and related notions such as **cognitive dysfunction, design plan**, and the like, deeply problematic? What is it for a natural organism—a tree, for example, or a horse—to be in good working order, to be functioning properly? Isn't "working properly" relative to our aims and interests? A cow is functioning properly when she gives the appropriate kind and amount of milk; a garden patch is as it ought to be when it displays a luxuriant preponderance of the sorts of vegetation we propose to promote. But here it seems patent that what constitutes proper functioning depends upon our aims and interests. So far as nature herself goes, isn't a fish decomposing in a hill of corn functioning just as properly, just as excellently, as one happily swimming about chasing minnows? But then what could be meant by speaking of "proper functioning" with respect to our cognitive faculties? A chunk of reality—an organism, a part of an organism, an ecosystem, a garden patch—"functions properly", it might be said, only with respect to a sort of grid *we* impose on nature, a grid that incorporates our aims and desires.

Now from a *theistic* point of view—a point of view that I accept—the idea of proper functioning is no more problematic than, say that of a Boeing 747's working properly. Something we have constructed—a heating system, a rope, a linear accelerator—is functioning properly when it is functioning in the way in which it was designed to function. But according to theism, human beings, like ropes and linear accelerators, have been designed; they have been designed and created by God. According to the theistic way of looking at the matter, we human beings have been created by God, and created in his own image; in certain important respect we resemble him. God, furthermore, is an actor, a creator, one who chooses cer-

tain ends and takes action to accomplish them. God is therefore a *practical* being. But he is also an *intellectual* or *intellecting* being. He holds beliefs, (even if there are significant differences between his way of holding a belief and ours[30]), has knowledge, apprehends concepts. In setting out to create human beings in his image, then, God set out to create them in such a way that they could reflect his capacity to grasp concepts and hold beliefs. Furthermore, he proposed to create them in such a way that they can reflect his ability to hold *true* beliefs. He therefore created us with cognitive faculties or powers designed to enable us to achieve true beliefs with respect to a wide variety of propositions—propositions about our immediate environment, about the past, about our own interior life, about the thoughts and feelings of other persons, about our universe at large, about right and wrong, about the whole realm of *abstracta* (properties, propositions, states of affairs, numbers, and the like) about modality, and about himself.

From this perspective it is easy enough to say what it is for our faculties to be working properly; they are working properly when they are working in the way they were intended to work by the being who designed them. This, I take it, is the basic root of the idea of proper functioning: an object is functioning properly if and only if it is functioning in the way it was designed to function. Of course from a theistic perspective we must append a qualification to the idea that positive epistemic status is a matter of proper functioning; for clearly enough **S**'s faculties could be working properly, in that root sense, even if **S**'s beliefs have no positive epistemic status for him. **S**'s faculties could have been designed by a mischievous Cartesian demon who finds it diverting to contemplate creatures who are both systematically deceived and proud of their status as the epistemic lords of the universe; or perhaps **S**'s epistemic faculties have been redesigned by an Alpha Centaurian scientist who cares nothing for their relation to truth. So even if we say that **S**'s faculties are working properly, it doesn't follow that **S** has positive epistemic status for his beliefs. What we must add is what the theistic view does add: that our faculties have been designed by a being who wishes to enable us to achieve a substantial degree of truth in a substantial portion of the range in which we form beliefs.

But can a nontheist also make use of this notion of proper function in understanding positive epistemic status? Well, why not? Can't anyone, theist or not, see that a horse, say, in suffering from a disease,

is displaying a pathological condition? Can't anyone see that an injured bird has a wing that isn't working properly? Or that an arthritic hand does not function properly? The question is whether theism is entailed by the claim that we have faculties that function properly or improperly (faculties that have a design plan). The question is whether, on the proposed account of positive epistemic status, the proposition that some belief has positive epistemic status for you entails the truth of theism. Of course it might trivially entail theism, by virtue of the fact (as the theist sees it) that theism is a necessary truth; alternatively, perhaps theism is not a necessary truth, but (as the theist may also think) it is necessary that all contingent beings have been created directly or indirectly by God. But of course if this is true then any account whatever of knowledge (or of anything else) will entail theism. The real question here is something different—something not at all easy to state clearly. Perhaps we can put it like this: the real question is whether the notion of proper function is linked with theism in such a way that the proposition **some organ or system of some organism is functioning properly**, entails theism by way of a series of steps each of which is obvious. Or perhaps the real question is whether there is a satisfactory nontheistic *analysis* or explanation of the notion of proper function. It is certainly not obvious either that there is any such entailment or that there is no such explanation or analysis. More vaguely, the question is whether the notion of proper function can be properly understood from a nontheistic perspective. But even if it can't, that is no real objection to the present account. We all have and constantly use the idea of the proper function of our cognitive faculties; we all have and use the idea of the dysfunction of various systems and organs of human beings and other organisms. If there is no adequate nontheistic way to understand this family of notions, then there lurks here, not an objection to the above account of positive epistemic status, but a powerful theistic argument.

And even if there is no good nontheistic analysis of **proper function**, the nontheist can accept something *like* this notion. Even if he doesn't think we human beings have been designed and created by a powerful and highly competent being who proposed to endow us with the ability to achieve true beliefs, he may nonetheless think of this idea as a convenient and useful fiction. He may join Hans Vaihinger in *Der Philosophie von Als Ob*, and explain proper functioning in terms of this fiction, as he sees it; he may say that our

faculties are working properly when they are working the way they would work if the theistic story were true.[31] He may therefore treat this story the way corresponding stories are treated by some who accept ideal observer theories in ethics, or social contract theories in political philosophy, or Piercian theories of truth, or possible worlds theories in metaphysics. I can sensibly explain what it is for an action to be right in terms of what an ideal observer would approve without adding that in fact there exists an ideal observer. A person can see possible worlds theory as a source of insight and understanding, even if he thinks it is not to be taken seriously as sober metaphysics. In the same way a nontheist could help himself to the theistic explanation of positive epistemic status, even if he thinks the notion of proper function has no very good nontheistic explanation.

There is a similar but slightly different tack he may take: He may take towards the idea of design the same attitude Bas van Fraassen takes towards possible worlds, modality, and unobservables in science. These, says van Fraassen, are pictures that guide our inference, but they are not to be taken seriously as part of the sober metaphysical truth of the matter. "Such fictions," he says, "are useful in giving an account of the surface phenomena—and there is in reality, nothing below the surface. In our case the phenomena are the inferential relations among statements, attested in the inferential behavior of those engaged in such discourse."[32] Each of these stances, admittedly, is perhaps vaguely uneasy or a bit awkward; but there is nothing initially incoherent about them. And in the same sort of way a nontheist can accept the present account of positive epistemic status, even if he thinks there is no good way to understand proper function and allied notions from a nontheistic perspective.

By way of conclusion then: the main contemporary accounts of positive epistemic status are all deeply flawed. Each founders on the same rock: each neglects to take into account the ways in which our cognitive faculties can fail to function properly, and each overlooks the crucial connection between positive epistemic status, on the one hand, and our cognitive faculties' functioning properly on the other. Indeed, positive epistemic status, as I see it, just is the proper functioning of our epistemic equipment. In a nutshell, a belief **B** has positive epistemic status for **S** if and only if that belief is produced in **S** by his epistemic faculties' working properly; and **B** has more positive epistemic status than **B*** for **S** iff **B** has positive epistemic

status for **S** and either **B*** does not or else **S** is more strongly inclined to believe **B** than **B***. This picture of positive epistemic status, obviously enough, needs articulation, development and qualification; nevertheless, I think, it is a better picture than any of its rivals.

Notes

1. And until Edmund Gettier showed us the error of our ways: see his classic "Is Justified True Belief Knowledge?" *Analysis* 23 (1963) pp. 121-123.
2. See , e.g., *Theory of Knowledge* (New York: Prentice Hall, 2nd Edition), p. 5.
3. See my "Reason and Belief in God" in *Faith and Rationality: Reason and Belief in God*, ed. A. Plantinga and N. Wolterstorff (South Bend: University of Notre Dame Press, 1985), pp. 24 ff.
4. Call the first belief 'A' and the second 'B'; there is the degree of positive epistemic status had by those beliefs that have no more of that quantity (for me) than B, the degree displayed by those that have more than B but less than A, and the degree enjoyed by those that have as much or more than A.
5. For a fuller and more complete version of the ideas of this section, see my "Chisholmian Internalism" in *Philosophical Analysis: a Defense by Example*, ed. David Austin (forthcoming).
6. Although this has perhaps been the dominant internalist tradition in modern (Western) epistemology, it is not the only one. To characterize the other one, we must back up just a bit. According to the second tradition, positive epistemic status is to be understood in terms of rational action. An action I take is rational in case it is appropriately connected with the attainment of my aims or goals—perhaps it is the action (of all those open to me) that is *in fact* most suited to achieving my goals, or perhaps it is the one that I *think* is the one best suited, or perhaps the one such that upon sufficient reflection I *would* think it most suitable. But then epistemic rationality is a special case of rationality generally, and is to be thought of in terms of the aptness of my beliefs or my belief forming policies to fulfill my epistemic goals. See Richard Foley's book *The Theory of Epistemic Rationality* (Cambridge: Harvard University Press, 1987) for an excellent contemporary development of a non-Cartesian variety of internalism.
7. See "Chisholmian Internalism", first section.
8. I ignore, here, the complication provided by the fact that in *Foundations of Knowing* Chisholm recasts his epistemological theory in terms of direct attribution of properties, thus abandoning the earlier formulation in terms of believing or accepting propositions. Everything I say about Chisholm's views can be restated so as to accommodate this shift, though in some cases at the cost of considerable complication.
9. But even if epistemic obligation or requirement is *sui generis*, it shares important elements of structure with moral obligation: there is super-

venience, defeasibility, the application of the *prima facie*/all things considered distinction, the characteristic relations among permission, obligation, and prohibition, and so on.

10. See "Chisholmian Internalism", first section.
11. Much of Chisholm's epistemological work is given over to the project of formulating and defending epistemic principles (see the earlier pages of this section). I do not propose to argue that these principles are mistaken; I mean to argue only that his suggestion as to what positive epistemic status *is*, is mistaken.
12. For example: suppose you were nonculpably convinced that the most likely way to achieve truth was to avoid thinking things over and to believe the first thing that came into your head; would it then be your duty to form beliefs in that way? Suppose you thought, after careful and dutiful reflection, that a given proposition was false: could it nonetheless be your duty to believe it? If you failed to believe it, could you properly be accused of dereliction of duty? See my "Reason and Belief in God", pp.35-37.
13. And it may well be less than that; see "Chisholmian Internalism".
14. For some suggestions as to how coherentism is related to foundationalism, see my "Coherentism and the Evidentialist Objection to Theistic Belief" in *Rationality, Religious Belief and Moral Commitment*, ed. R. Audi and W. Wainwright, (Ithaca: Cornell University Press, 1986).
15. An outstanding exception is Keith Lehrer's book *Knowledge* (Oxford: Oxford University Press, 1974) surely as full and articulate development of coherentist thought as is presently to be found. Although I don't have the space here to give a critical account of Lehrer's conception of coherence (but see my "Coherentism and the Evidentialist Objection to Belief in God") I intend the criticisms that follow to apply to it. Of course there are also probabilistic or Baysian coherentists; I don't here have the space to discuss this position.
16. See, for example, Fred Dretske's *Knowledge and the Flow of Information* (Cambridge: The MIT Press, 1981); Alvin Goldman's "What is Justified Belief" in *Justification and Knowledge: New Studies in Epistemology*, ed. G. Pappas, (Boston, Dordrecht, London: D. Reidel); Robert Nozick's *Philosophical Explanations* (Cambridge: Harvard University Press, 1981; and Marshall Swain's *Reasons and Knowledge* (Ithaca: Cornell University Press, 1981).
17. Nozick, *Philosophical Explanations*, pp.172-178. Nozick points out that conditions (1)-(4), as they stand, are neither necessary nor sufficient for knowledge; he tries to mend matters by adding an epicycle involving "methods" of believing or coming to believe: "This leads us to put forth the following position: **S** knows that *p* if there is some method via which **S** believes that *p* which satisfies conditions (1)-(4), and that method is not outweighed by any other method(s) via which **S** actually believes that *p* that fail to satisfy conditions (3) and (4)". The idea of coming to believe via a "method" isn't very clear; insofar as I see what Nozick means here, I have tried to construct the counterexamples I propose below in such a way that they involve **S**'s coming to believe via just

one method, so that this complication will not be relevant.

18. Hence on Nozick's view, I don't know (4), do know (3), don't know the disjunction of (3) with (4) (if it were false, (4) would be false, so that I would still believe the disjunction) but do know the conjunction of (3) with (4)!
19. Cambridge: MIT Press, 1981.
20. See Richard Foley "Dretske's 'Information Theoretic' Account of Knowledge" forthcoming in *Synthese*.
21. Thus according to Reichenbach, the appropriate reference class is the narrowest class such that we know that Paul is a member of it, and such that we have reliable statistics for the frequency of the attribute in that class. Others have made other suggestions; Wesley Salmon, for example, suggests that the appropriate reference class is the broadest homogenous class containing the instance in question. (See his *Foundations of Scientific Inference*, (Pittsburgh: The University of Pittsburgh Press, 1966) p. 91. All the suggestions made, however, result in the 'probability' for a singular proposition's being relative to some body of knowledge.
22. *The Behavioral and Brain Sciences* (1983) p. 83.
23. Objection: in this example (as in the next) I am conforming to the *letter* of D_5 but not to the spirit of the underlying intuition. That intuition is that in the relevant epistemic situations, there is a source **s** that sends a signal **r** having the property **G**; **K** knows that **s** if **F** in case **r's being G** (1) causes **K** to believe that **s** is **F** and (2) carries the information that **s** is **F** (carries that information with respect to **K**). But my last two examples both work by *identifying* **r's being G** with **s's being F**.
Reply: In these objections I did indeed collapse **r's being G** into **s's being F**. (My reason for so doing was only to avoid avoidable problems about whether **P(s's being F | r's being G)** really equals 1.) But that is an inessential feature of these examples; we can easily amend them to satisfy the objection. Accordingly, revise the present case as follows: I implant a source of high energy radiation in your dog Spot; it is a law-like truth that any dog in which a source of high energy radiation has been implanted will lose its hair within seven days; Spot's emitting this high energy radiation causes a brain lesion in you which in turn causes you to form a large number of wildly false beliefs about Spot (that he is in fact a mermaid, that he can speak French but refuses to out of sheer obstinacy, etc.) but also causes you to form the true belief that Spot will lose his hair within the next two weeks. You have no evidence of any sort for your belief and much evidence against it. (You have just had Spot examined by a team of tonsorial experts who assure you that he is entirely normal along these lines.) Here **r's being G** is not collapsed into **s's being F** (**r's being G** is **Spot's emitting high energy radiation** and **s's being F** is **Spot's losing his hair within two weeks**). You satisfy the conditions laid down by D_5 for knowledge; but surely you don't know. The next example (about the doubly rigged lottery) can be similarly amended.

24. "What is Justified Belief?" in *Justification and Knowledge: New Studies in Epistemology*, ed. George Pappas (Dordrecht: D. Reidel, 1979), p. 10.
25. *Monist*, 1986.
26. And perhaps the "such and such lightings and atmospheric conditions" should be deleted; perhaps Goldman's idea is that the types in question must be stateable in psychological terms alone.
27. Robert Shope has called my attention to a passage where Goldman may be proposing a further restriction on the sorts of processes that are relevant: "Justifiedness seems to be a function of how a cognizer deals with his environmental input, i.e., with the goodness or badness of the operations that register and transform the stimulation that reaches him" (12-13). Shope suggests (private communication) that perhaps Goldman intends the processes under discussion to be ones that transform or operate upon external stimuli, a condition that is not met in the brain tumor case. Very well, alter the case accordingly; let the relevant process associated with the tumor be one that takes any visual stimulus, say, as input and yields as output the pathological belief in question.
28. See below, section IV B.
29. See above section III C.
30. See my "Justification and Theism", *Faith and Philosophy*, Oct. 1987, last section.
31. But here he must be careful. If theism were true, then presumably there would be something like Calvin's Sensus Divinitatis, a manysided disposition to accept belief in God or propositions about God in a wide variety of circumstances. (See my "Reason and Belief in God", p. 80-82.) But then our faculties would be functioning properly when we form such beliefs in the basic way—that is, immediately, without believing on the evidential basis of other propositions. And then such belief in God would have positive epistemic status—a conclusion the nontheist may wish to avoid.
32. "Probabilities of Conditionals" in *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science, I*, ed. Harper and Hooker, (Dordrecht: D. Reidel Publishing Co., 1976) p. 267.