PHILOSOPHICAL TOPICS
VOL. 23 NO. 1, SPRING 1995

How to Think about Reliability

William P. Alston
Syracuse University

I

Theories of epistemic justification and of knowledge in terms of *reliability* have been prominent lately. The way it typically goes is that justification of belief is explained in terms of reliability and then having a justified belief that p is taken as necessary for knowledge that p. Among those who propound theories of this sort are Alvin Goldman,[1] Marshall Swain,[2] and Frederick Schmitt.[3] The basic intuitive idea of a reliabilist theory of justification is that justified belief is belief that results from a reliable belief-forming process.[4]

> If S's believing p at t results from a reliable cognitive belief-forming process . . . then S's belief in p at t is justified.[5]

In this paper I will not be concerned with setting out, defending, or attacking a reliabilist account of justification or knowledge. Indeed, I will not be directly concerned with epistemological issues at all. What I will treat is the *concept* of *reliable belief formation*. I am prompted to do this by the fact that among the many criticisms that reliabilist epistemology has evoked, a prominent place is occupied by conceptual criticisms. It has been claimed more than once that the notion of *a belief's being formed by a reliable process* is too indeterminate to permit any satisfactory and objective answer to

the question of whether it applies in a given case. If that is so, the theory can't get off the ground. If our grasp of *reliability* is such that we are, in principle, unable to tell whether a given belief is formed reliably or not, we cannot take the first step toward considering whether, for example, a belief is justified *iff* it is reliably formed. Reliabilist epistemologists, including Goldman and Schmitt, have taken such criticism very seriously and have attempted to respond to them. I find these attempts not to be very successful.[6] In this essay I hope to do better. My aim is to articulate a conception of *reliable belief formation* that can be used to set out a reliabilist theory of epistemic justification that is at least not internally incoherent.

II

The criticism on which I will concentrate has to do with what Richard Feldman has called the **Problem of Generality**.[7] The problem arises as follows. A particular belief is generated by a particular psychological process. But a particular process, with a certain spatiotemporal location, is not the sort of thing that can be more or less reliable. Reliability or the reverse attaches only to what is repeatable, to what has, actually or potentially, a number of instances. At least this is true if we are thinking of events or processes, as reliabilists typically are. It is more common in ordinary speech to attribute reliability and unreliability to mechanisms, like thermometers, carburetors, and clocks, or to medicines, or to sources of information, whether persons or books or newspapers. These are all *substances*, *things* of one kind or another, and here *repeatability* and *instances* are not in the picture. Nevertheless, an analogous point can be made. Though a clock or an encyclopedia or a medicine is not "repeatable," it is something that can be operated, consulted, or used many times, and so there is something like repeatability here, viz., repeated employments. So the general point is that to be assessable as reliable or the reverse, something must, actually or potentially, provide a range of cases of the appropriate sort. For reliability is always a matter of the incidence of favorable outcomes in a multitude of instances or employments of the item in question. And if there is not at least the possibility of such a multitude, the question of the proportion of favorable cases cannot arise. What counts as a *favorable* case differs, of course, from one application to another. With medicines, it is a certain kind of medical improvement. With clocks, it is registering the correct time. With maps, it is accurate representation of location and distance. For the case at hand—belief-forming processes—it is truth. A reliable belief-forming process is one that usually yields *true* beliefs. And, to get back to the original point, a particular process that takes place at a particular precise time is not the sort of thing

that does or does not enjoy a favorable ratio of true beliefs among its products. It occurs just once; the one belief it produces is either true or false, and there's an end to it. Hence, as is regularly said by both friend and foe, it is a *type* of cognitive process, rather than a particular process (a *token*), that can be assessed for reliability. Consider a particular case of *inferring a generalization from a number of instances*. If we want to ask how reliable that process is, we have to consider the *type* of inference to which it belongs and seek to determine the proportion of true generalizations among the (actual or possible) products of *tokens of that type*.

Before continuing, I had best make it explicit that reliability is a degree concept. An instrument, book, or person can be more or less reliable. Thus when one uses 'reliable' and 'unreliable' as absolute, yes-or-no terms, one must, if one is talking sense, have some minimum degree of reliability in mind as a necessary condition for speaking of the item as reliable tout court. We don't want to be too rigid about this. One's conception of reliability, like many other concepts, might be less than fully precise without lacking content. Being reliable might be thought of as "usually" or "most often" exhibiting favorable cases, or as there being a "high proportion" of favorable cases, or as "most" of the cases being favorable, rather than being thought of in terms of a precise minimum percentage.[8] I mention this point only to set it aside. Though it is an important point for the development of a reliabilist epistemology, it does not bear on the issues I will be discussing in this essay.

These are the rudiments of the situation. Thus far the difficulty has not shown itself. But it is not far below the surface. If we seek to determine whether a particular belief is justified by determining whether it was produced by a reliable process, the above remarks indicate that we must determine whether the *type* of belief-forming process of which this belief formation is an instance usually does or would generate true beliefs. And the difficulty is that there is no unique type that meets that description. Any particular process is an instance of indefinitely many process types. I will let Feldman make the point.

> [T]he specific process token that leads to any belief will always be an instance of many process types. For example, the process token leading to my current belief that it is sunny today is an instance of all the following types: the perceptual process, the visual process, processes that occur on Wednesday, processes that lead to true beliefs, etc. Note that these process types are not equally reliable. Obviously, then, one of these types must be the one whose reliability is relevant to the assessment of my belief. Intuitively, it seems clear that the general reliability of processes that occur on Wednesday or processes that lead to true beliefs is not relevant to the assessment of my belief. The reliability of the visual process or of the perceptual process may well be important.

> Let us say, then, for each belief-forming process token there is some "relevant" type such that it is the reliability of that type which determines the justifiability of the belief produced by that token. Thus, the reliability theory can be formulated as follows:
>
> > (RT) S's belief that *p* is justified if and only if the process leading to S's belief that *p* is a process token whose relevant process type is reliable.
>
> In order to evaluate (RT), we need some account of what the relevant types of belief-forming processes are. Without such an account, we simply have no idea what consequences the proposal has since we have no idea which process types are relevant to the evaluation of any particular beliefs.[9]

Here is another formulation of the same complaint by Alvin Plantinga.

> 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?[10]

Thus there is no unique type to which a particular process token belongs. Each token belongs to indefinitely many types. Hence to construct a reliabilist theory of justification, we must find some way to pick from this embarrassment of riches a type on the basis of which a judgment of the reliability of this particular belief formation can be made. And, they continue, there would seem to be no satisfactory way of doing this. Here is Feldman's picture of the situation.

> In coming up with an account of relevant types, defenders of the reliability theory must be guided by the following point. If relevant types are characterized very narrowly then the relevant type

> for some or all process tokens will have only one instance (namely, that token itself). If that token leads to a true belief, then its relevant type is completely reliable, and according to (RT), the belief it produces is justified. If that token leads to a false belief, then its relevant type is completely unreliable, and, according to (RT), the belief it produces is unjustified. This is plainly unacceptable, and in the extreme case, where every relevant type has only one instance, (RT) has the absurd consequence that all true beliefs are justified and all false beliefs are unjustified. We can say that characterizing relevant types too narrowly leads to the "The Single Case Problem."
>
> A very broad account of relevant types of belief-forming processes leads to what we may call "The No-Distinction Problem." This arises when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type. For example, if the relevant type for every case of inferring were the type "inferring," then (RT) would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type.
>
> The problem for defenders of the reliability theory, then, is to provide an account of relevant types that is broad enough to avoid The Single Case Problem but not so broad as to encounter The No-Distinction Problem. Let us call the problem of finding such an account "The Problem of Generality."[11]

Feldman then goes on to argue at some length that attempts to solve this problem have not been successful, and he suggests that the prospects for doing better are poor.[12]

III

Now if I were to accept Feldman's challenge on his terms and attempt to respond, I would be engaged in what I said I would not do in this essay, viz., enter into the merits and demerits of reliabilist theories of justification. For to find a way of specifying "relevant types" that avoids both his Scylla and his Charybdis is to find a way of assigning tokens to types that makes the former line up with intuitive judgments of justification in the right way, and I have disavowed any intention of doing that here. What I will do instead is to challenge the basic assumption that lies behind the challenge, viz., that there are no objective, psychological facts of the matter that pick out a unique type as the one of which a particular process is a token. It is only if that assumption is true that the reliabilist epistemologist is faced with the daunting task of providing a principled way of selecting a single winner from this plethora of candidates. And note that even if she should succeed in that task, she would

still be subject to the complaint that the assignment of each token to a unique type has been rigged to fit an antecedent decision as to the epistemic status of the belief, thereby giving rise to the suspicion that reliability is not the most basic determinant of justification after all. Hence a great deal is riding on the supposition that there are no objective, *nonepistemic* facts that determine a unique type assignment for each token belief formation. It is that supposition I will seek to discredit.

Before turning to that, however, I can use purely conceptual, nonepistemic, considerations to show how to avoid Feldman's Scylla, the "Single Case Problem." That can be disposed of just by making it explicit that reliability is *not* a matter of actual track record but rather is a "propensity" or "disposition" notion. To say that a thermometer, medicine, or atlas is *reliable* is not to make a report of the relative frequency of favorable outcomes in the cases in which it has been used up to this point. It may never have been used at all yet, but that doesn't keep it from being reliable or unreliable. We may not be able to tell how reliable it is if it hasn't yet been used, but here as elsewhere it is a great mistake to conflate *X*'s *being P* with our ascertaining, or being able to ascertain, that it is *P*. A thermometer may be perfectly reliable, as we may discover in time, even though it has just rolled off the assembly line and has not yet been put to the test. Nor is the reliability of a device a function of the proportion of favorable outcomes of uses over its entire life history—past, present, and future. An atlas might be perfectly reliable even though no one ever opens it or consults it for any purpose. In this respect reliability functions like other dispositional properties. A rubber band can be elastic even though it is never stretched and never has the chance to manifest that disposition. The applicability of a dispositional term depends on whether the appropriate manifestations *would* result from the satisfaction of the relevant antecedent conditions in a suitable range of cases, whether or not such a spread of cases, or any cases at all, are ever forthcoming. An elastic substance is one that *would* resume its shape *if* deformed. A responsible person is one who *would* fulfill whatever commitments he or she made, or most of them. A reliable instrument is one that *would* usually deliver favorable results over an appropriate range of cases *if and when* they occur. The same point holds, with appropriate adjustments, for process types. A type of belief-forming process is a reliable type provided that tokens of that type *would* generate true beliefs in a large proportion of a suitable run of cases.

In denying that reliability *is* a matter of actual frequencies, I am not denying that the most direct way of assessing a device or a process type for reliability is to ascertain the frequency of favorable outcomes in a suitable range of cases. That is not the only way, but it is certainly the most direct way. And if we couldn't use that approach for many dispositional properties, we would not be in a position to employ less direct procedures that require

already knowing how to apply other dispositional concepts. Thus, having ascertained by actual-frequency counts that certain human perceptual belief-forming processes are reliable, we thereby have some basis for inferring that processes sufficiently similar to those are likely to be reliable also. But the latter inference requires that we have already done some frequency count spadework for some processes.

Note that I have said that we need to *make explicit* the point that reliability *is* a dispositional or propensity notion, rather than a track-record notion. I did not say that we need to *develop* the notion in this direction or *modify* it so that it takes this shape. It seems perfectly clear to me that as we ordinarily use the term 'reliable' it functions in the way I have just described. Feldman presents the propensity construal as something a reliability theorist might be driven to in order to answer objections.[13] And even Goldman is much too permissive on this point.

> I have characterized justification-conferring processes as ones that have a 'tendency' to produce beliefs that are true rather than false. The term 'tendency' could refer either to *actual* long-run frequency, or to a 'propensity', i.e., outcomes that would occur in merely possible realizations of the process. Which of these is intended? Unfortunately, I think our ordinary conception of justifiedness is vague on this dimension. . . .[14]

But I don't see how anyone could possibly think that 'tendency' might mean "actual long-run frequency." Nor do I see how anyone could reasonably suppose that 'reliable' (that's the term we should be discussing rather than 'justified') is vague on the distinction between propensity and actual frequency. That's certainly not my sense of the way these terms work.

Before continuing the main line of argument, I had better make explicit an important distinction between two types of belief-forming processes. In "What Is Justified Belief?" Goldman distinguishes between "belief-independent" and "belief-dependent" processes. The former take no doxastic inputs; the latter take at least some. Goldman characterizes reliability differently for the two types, and I will follow him in this. The formulation just given applies straightforwardly to belief-independent processes. In their case, reliability is a simple matter of the proportion of true outputs that would result from an appropriate range of cases. But the latter division is more complicated. There we have to give separate attention to the transfer of epistemic status across the inference and to the status that is being transferred. Goldman handles this by treating separately the reliability of the process and the reliable formation of the output belief.

> A belief-dependent process is *conditionally reliable* when a sufficient proportion of its output-beliefs are true *given that its input-beliefs are true.*

> 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.[15]

Dividing the situation up in this way enables us to assess the process for reliability without worrying about the epistemic status of the inputs and then to use that result, along with the epistemic status of the inputs, to give conditions for the belief to satisfy reliabilist conditions for justification. With belief-independent processes there is no need for this two-stage treatment, since the inputs (experiences) are not the sort of things that have epistemic status.

The application of my point that reliability is a propensity notion to Feldman's worries about "the single case problem" is quite straightforward. Since reliability doesn't hang on actual frequency, there is no excuse for supposing that if a belief-forming process type is exemplified only once and produces a true belief it is perfectly reliable, whereas if it produces a false belief on its only instantiation it is completely unreliable. Its place on the reliability dimension depends on what the proportion of true beliefs *would be* in a suitable range of instantiations, not on what actually results when it is realized.

In contrasting a track-record construal with a propensity construal, I have presented the latter in terms of the frequency that would be displayed by a "suitable" or "appropriate" range of cases. What does it take for a range of cases to be "suitable"? First, they must be sufficiently numerous. What counts as sufficient will vary with different subject matters. We learn from experience how much homogeneity or heterogeneity to expect in a number of cases of a given sort. The more variation we have learned to expect across instances, the larger the sample we need. Since people vary much more than chemical substances, we need a larger sample for opinion or attitude research than for determining atomic weight. Second, the cases must be sufficiently varied along relevant dimensions to rule out, so far as possible, the hypothesis that the results are due to factors other than the characteristic being tested for. If we are testing for arithmetic ability, we wouldn't want to confine ourselves to very easy or very difficult problems; we would want to include a spread of difficulty in the test. Similarly, if we are testing a perceptual belief-forming process for reliability, we would want to vary the cases with respect to such factors as distracting "noise" from other psychological inputs and the degree of discrimination required between inputs close to each other in time and/or space. These two points are applications to this topic of criteria that are applied to any statistical or sampling procedure.

But there is another factor that must be taken into account here. What the last paragraph suggests is that a belief-forming process will be deemed

reliable provided its exercise in a suitable number and range of cases would result in a high proportion of true beliefs. But this would seem to be false for many of the processes we would ordinarily regard as reliable, so long as no restrictions are put on the circumstances in which the process occurs. Many human belief-forming processes would exhibit sharply different degrees of reliability over the whole range of conceivable situations. This may not be true of all. Some have supposed that introspection must yield only truths, no matter what, and some have taken the apprehension of propositions as self-evident to enjoy a like immunity to error. Moreover, if we restrict ourselves, as I have been doing, to assessing inferential processes in terms of *conditional* reliability, then if we ignore Cartesian doubts about the a priori, it seems that in no possible situation would deductively valid inferential processes fail to be conditionally reliable. But look at perception and nondeductive reasoning. For any perceptual mechanism that produces mostly truths in the situations in which it is actually exercised, there are possible situations in which that reliability would be sharply reduced. And not just logically possible situations either. Even if the Cartesian demon is only logically possible, it is well within our powers to arrange environments in which a normal person, utilizing normal mechanisms of perceptual belief formation that serve us well in run-of-the-mill situations, would, usually or always, be led astray. We need only manufacture realistic enough look-alikes, or do something more ambitious with holographs, or something still more ambitious with direct brain stimulation. If the range of cases in terms of which the process is assessed for reliability were restricted to situations like these, it would score very low, even if the score would be high when tested in more familiar situations. To illustrate the problem with respect to nondeductive reasoning, we must have recourse to more recherché possibilities. It does seem at least logically possible that the world should be such that the modes of nondeductive reasoning in which we are most confident would fail to produce mostly true beliefs from true premises and so would score low on conditional reliability. Consider induction by simple enumeration. Surely there are possible worlds in which when there is a high proportion of F's in G's in what we take to be a properly constituted sample, it is usually the case that most G's overall are not F. If all else fails, we can introduce a Cartesian demon that delights in arranging things so that such inferences are usually, or invariably, frustrated.

Thus we are faced with a question as to the range of situations in which a process must yield mostly true beliefs in order that it qualify as reliable. If we require reliability over all possible situations, we will, at best, be left with a sharply reduced set of reliable human belief-forming habits. And if not that, how shall we demarcate those situations over which the test sample must range in order to give a relevant result?

I can't think of any better answer to this question than the following. The requirement for reliability is that the process would yield a high proportion of truths over a wide range of situations *of the sort we typically encounter*. Obviously, this is far from precise. It doesn't draw a sharp boundary between typical and atypical. Moreover, it leaves open the possibility that the boundary, such as it is, can shift over time. What was atypical up to now may become typical with cultural, technological, or other changes. However, I believe that this suggestion has the right kind and degree of sloppiness for the concept of reliability we want for epistemic purposes. It does unequivocally rule out clearly atypical situations—Cartesian demons, brains in vats, and the like.[16] And it makes a judgment of reliability dependent on our actual situation as human beings in the environments in which we actually find ourselves. This is what we need to capture the intuitive notion of reliability that is involved in reliabilist epistemologies. If I claim that my thermometer is reliable, it is no refutation to point out that it would not give an accurate reading on the sun. Similarly, if I claim to be able to accurately determine, by vision, when I am standing in front of a beech tree, it would be no refutation of that claim to point out that I could not do this if I were receiving direct stimulation of the visual cortex in a physiological laboratory. When I make a judgment of reliability—whether for an instrument, a documentary source, a psychological mechanism, or whatever—I have in mind, at least implicitly, a range of situations with respect to which the claim is being made. What happens outside that range is simply irrelevant to the claim.

IV

Now I can return to my central task—challenging the assumption that there are no nonepistemic facts that determine a unique type to which a given belief-forming process belongs. Of course, in a way I agree that a given process token does not belong to only one type. With a process token, as with any other particular, any of its properties can be said to be correlated with a type to which it belongs; and its properties are indefinitely numerous. Nevertheless, some types, in this maximally generous logical sense, are ontologically rooted, fundamental, or important in ways indefinitely many others are not. Even if it is true that you and I belong to indefinitely many classes, such as *objects weighing more than ten pounds*, *objects that exist in the twentieth century*, *objects mentioned in this paper*, etc., etc., it is still the case that membership in the class of human beings is fundamental for *what we are* in a way those others are not, just because it is the *natural kind* to which we belong. I shall suggest that something analogous is true of belief-forming

processes—that there are fundamental considerations that mark out, for each such process token, a type that is something like its "natural kind," and hence that in thinking of belief-forming process types, we are not awash in a sea of indeterminacy as Feldman and company suppose.

To carry this through we will have to think of belief-forming processes in a somewhat different way than is common with both friends and foes of reliabilism. First, we will have to draw boundaries around the process more narrowly than is often done. Most discussions of the current issue are not very specific about the extent of the processes under discussion. One crucial question concerns whether we should think of the process as including everything in the causal ancestry of the belief or only some segment thereof. In particular, should we include events outside the organism or limit them to certain intra-organismic, or intra-psychic, events that function as proximate causes of the belief? Critics typically follow the lead of some reliabilists in leaving this hazy. As an example of the latter, Armstrong's version of reliabilism runs something like this. *S*'s belief that *p* is reliably formed provided that *S* has some property, *H*, such that it is nomologically necessary that if a subject that is *H* forms a belief that *p*, that belief is true.[17] This formulation is not explicitly in terms of processes, but it could be rewritten in those terms, as follows. *S*'s belief that *p* is reliably formed provided the process that generated that belief has some property, *H*, such that it is nomologically necessary that any belief that is generated by a process with that property is true. This obviously puts no restriction on the extent of the process. Any characterization of the process that will yield the nomological necessity in question will ensure reliability of belief formation.[18] As for the critics, we find Feldman and Pollock considering perceptual processes that include factors outside the subject, such as distance from the object and lighting conditions.

However our paradigm reliabilist, Goldman, is more specific.

> In addition to the problem of 'generality' . . . there is the problem of the '*extent*' of belief-forming processes. Clearly, the causal ancestry of beliefs often includes events outside the organism. Are such events to be included among the 'inputs' of belief-forming processes? Or should we restrict the extent of belief-forming processes to '*cognitive*' events, i.e., events within the organism's nervous system? I shall choose the latter course, though with some hesitation. My general grounds for this decision are roughly as follows. 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. . . . A justified belief is, roughly speaking, one that results from cognitive operations that are, generally speaking, good or successful. But '*cognitive*' operations are most plausibly construed as operations of the cognitive faculties, i.e., 'information-processing' equipment *internal* to the organism.[19]

This seems to me just the right thing for a reliabilist to say on this point. If the epistemic status of a belief is a function of the reliability of the process that generates the belief, it is the reliability of the *psychological* process that is crucial. Looking at perceptual belief formation, no matter how exemplary the path of the light rays from the surface of the perceived object to the retina, and no matter how finely tuned the neural transformations involved in the pathway from the eye to the brain, if the belief is not formed on the basis of the conscious presentation (and/or its neural correlate) in a truth-conducive way, the belief will lack the epistemic desideratum that is stressed by reliabilism. It is that final step that is crucial. I note that by the time he came to write *Epistemology and Cognition* Goldman had lost the "hesitation" of which he speaks in the above passage. There he deals with *psychological* processes from the outset. See particularly chapter five.

But this limitation of belief-forming processes to the psyche does not significantly reduce the embarrassment of riches where types are concerned. A purely cognitive belief-forming process will also be of indefinitely many types, including such undesirable intruders as *happening on a Wednesday* and *generating a true belief.* Moreover it will belong to types of all levels of generality. If it is a visual belief-forming process, it will be of the type *forming a belief on the basis of such and such a kind of visual presentation, forming a belief about a tree, forming a belief about something in the vicinity, forming a belief on the basis of vision, forming a belief on the basis of perception*, and so on. We are still drowning in indeterminacy.

But decisive help is near. In fact, the germ of it is to be found in the very essay of Goldman's from which I have been quoting, though, as we shall see, he fails to take advantage of that idea to solve the generality problem.

> We need to say more about the notion of a belief-forming '*process*'. Let us mean by a 'process' a *functional operation* or procedure, i.e., 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.[20]

The crucial point here is that every belief formation involves the activation of a certain psychologically realized *function*. That activation yields a belief with a propositional content that is a certain function (the function the psychological realization of which is activated here) of the proximate input. The function involved will determine both what features of the input have a bearing on the belief output and what bearing they have, i.e., how the content of the belief is determined by those features.[21] In order to bring our talk of reliability of belief closer to such paradigm subjects of reliability attribution as thermometers, medicines, and atlases, let's say that a psychologically realized belief-formation function constitutes a psychological *mechanism*. If you don't like this terminology, either because it sounds too

"mechanistic" or because it threatens to populate the mind-brain with unmanageably many separate black boxes, we can just as well use other terms. We can think of the psychological realization of the function as a *habit* of forming a belief with a certain propositional content that is a certain function of certain features of the input or as a *disposition* to do that. Or if you prefer *act*-psychology, we can think of the subject's having the power to "take account" of certain features of inputs and, on the basis of that *taking-account*, form a belief with a content that is a certain function of those features. The common thread running through all this is that it is part of the constitution of the psyche to be so disposed that upon being presented with certain kinds of input a belief is generated with a content that is a certain function of certain features of that input. Remembering the variety of terminology available, I will in the ensuing oscillate between speaking of **habits** and of **mechanisms**.

Let's descend from these high levels of generality and look at some examples. Consider the formation of a visual perceptual belief that a maple tree is in front of one. The input will be a visual "presentation" of a certain sort, one that involves the perceived object's *looking* a certain way.[22] The mechanism that is activated will take account of certain phenomenal features of the presentation, while others will play no role. Certain shape features, certain color features, the spatial distribution of variously colored regions, and contrasts with the surrounding field will be "picked up" by the mechanism, while others will be ignored. As for the latter, many details of the presentation could have been different without changing the content of the belief generated. The tree could have looked larger or smaller, the bark could have looked rougher or smoother, and so on. It all depends on what function is operative. Obviously, if the function were one that delivered a belief about size or finely discriminated bark texture, the features that are irrelevant to whether one believes merely that there is a maple tree in front of one would have been relevant.

In this example I assumed that only the visual experience (presentation) functioned as an input. It is widely held that in every case of perceptual belief formation other beliefs of the subject play a role in shaping the doxastic output. Whether or not that is so, it is clear that in many cases the input is partly doxastic. As an example take an "individual recognition" case, rather than a "kind recognition" case like the previous one. Upon seeing a house I form the belief that it is your house. Actually there are a number of houses in the world that look just like yours, on a fleeting glance. But yours is the only one on this block of this street in this town that looks like yours. Hence, let us say, my identification of the house as yours (my formation of the belief that this is your house) is influenced by my knowledge (belief) that I am on this block of this street in this town, as well as by features of the visual presentation.

These examples may give the impression that the function involved in any perceptual belief formation is extremely specific, relating very detailed features of experiential input to a unique belief content. In view of the heavy weather made by Feldman and others about the problem of navigating between too specific and too general a relevant process-type, it will be pertinent for us to consider for a moment where psychologically realized belief-forming functions stand on that dimension. In a word, the answer is that the operative function can be of various degrees of generality. Sticking with perceptual belief formation for the moment, the function could be so specific as to take only precisely defined experiential features as input and issue only beliefs with a particular propositional content. That is the way I was thinking of the above examples. But there are other possibilities. Consider attributions of color to perceived objects. Here we have what is plausibly regarded as a single function that maps the position of certain aspects of a visual input on several color-relevant dimensions onto a belief that the object seen is of a certain color. Here we can be confident that the function is unitary because of the systematic character of the mapping. However I take this to be the exception rather than the rule for perceptual belief formation. For most propositional contents we don't have the possibility of systematic mapping that we have for color. Hence I will generally think of perceptual belief-forming functions as maximally specific.

Now let's look at a couple of inferential belief formations. It should be clear that on the approach I am suggesting the "functions" involved in deductive inference will be principles of inference. Think of a case of hypothetical syllogism. I reflect that if I refuse your request for a raise you will quit your job, and if you quit your job I will be unable to find a replacement in time for a big contract that is impending. I, naturally, infer that if I refuse your request for a raise I will be unable to find a replacement in time for that contract. The function that yields a belief that is related in that way to the input beliefs is a psychological realization of the principle of inference called 'hypothetical syllogism'. Obviously, we shouldn't suppose that only valid principles of deductive inference are psychologically realized. Some people, unfortunately, are so constituted as to regularly form beliefs related to the input beliefs in the pattern known as affirming the consequent. Such a function would yield the belief that *it rained last night* from the input beliefs that *if it rained last night the grass would be wet* and *the grass is wet.*

Nondeductive inference presents a more complicated picture. If I arrive at a generalization from knowledge of various instances, I typically take into account not only the instances in question but also facts about the subject matter that indicate what kind of sample I need in order to justifiably move to the generalization. As I pointed out above, some ranges of fact are more homogeneous than others and require less size and variety in the sample for a sound generalization. To be sure, we must keep reminding ourselves that

we are dealing here with what functions are actually psychologically realized, not only the ones that are reliable. But I think we may safely assume that in many cases considerations of the sort just mentioned figure in the input to an induction by simple enumeration. Again, consider inference to the best explanation. If I am trying to explain the presence of a pool of water on the floor of my basement, the relevant inputs will include not only beliefs about what the state of affairs is currently in my basement but also a list of possible causes of the water's being there, and considerations that bear on the likelihood of each of these causes having been operative.

The reader will, no doubt, have noticed that in the above I switched at a certain point from speaking of belief-forming *processes* to speaking of belief-forming *habits* (mechanisms, dispositions . . .). This may give the impression that my *activation-of-realized-functions* construal of belief formation is simply a different approach to the matter that has no relevance to the approach in terms of the reliability of *processes* of belief formation. But that is not the case. Habit and process are two aspects of the same phenomenon—belief formation. I am still speaking of belief-forming processes, but I have, following Goldman, made the conception of such processes more specific by thinking of such a process as the operation of a *habit* such as I have been describing. Whenever a belief-forming habit that involves a certain input-belief output function is in operation, there is a *process* of belief formation that consists of the input giving rise to the output in accordance with the function. It is just that instead of thinking, very unspecifically, of the process as whatever is involved in the causal history of the belief, I have limited the perspective to the proximate stage of that process and construed it in accordance with the input-function-output model.

V

The time has come to apply all this to the generality problem. The application is very simple; it has probably already leaped to the eye of the reader. **The function determines the relevant type**. I form the visual belief that a car is parked in front of my house. What type of belief-forming process is such that *its* reliability is crucial for the epistemic status of that belief, according to reliabilist epistemology? The type that is defined by the operative function, viz., *belief formations that proceed in accordance with the function that is involved here*. In other words, the particular process, by virtue of being a *functional* mapping of input features onto output content, has a *built in* generality that is provided by the function. The function *is* something inherently general, and it defines the type the reliability of which, according to reliabilism, is crucial for the justificatory status of the belief in question.

Let me say a bit more about just why the type defined by the operative function is the one to consider if we are interested, from the standpoint of reliabilist epistemology, in how reliably this particular belief was formed on this occasion. The type determined by the function has this special status just because it reflects or embodies the actual dynamics of the process, what is responsible for *this* belief with *this* content being formed on *this* basis. Hence if we assume, as reliabilist epistemology does, that the epistemic status of a belief is a function of its proximate causal history, of what led to its acquisition, then it is this type the reliability of which should be considered. To be sure, that assumption can be contested, but that issue lies outside the bounds of this paper, which is concerned with clearing up conceptual problems in reliabilist epistemology.

To return to the visual belief that a car is parked in front of my house, it would be an immense labor to spell out the function involved here in complete detail, since it would mean enumerating the features of the visual presentation that led to my identifying what I saw as a car parked in my driveway. (Specifying the output side of the function—the belief content—is no problem.) I can, of course, identify the experiential input accurately, though not analytically, by using the content of the belief output. I can say that the mechanism generated the belief that there is a car parked in my driveway on the basis of an input that consisted in *my being appeared to "car-parked-in-my-drivewayly,"* or in *its being just as if I were seeing a car parked in my driveway*, or in *an object's looking like a car parked in my driveway.* But whether we identify the perceptual input in analytically illuminating terms or in output-dependent terms, the basic point is the same. The type of process the reliability of which is relevant to the epistemic assessment of the belief is the one defined by the function, which is in turn defined by a certain way of going from input features to output features. The question of reliability that is of significance for the epistemic status of this belief is the question of how reliable *this* habit, the one defined by *this* function, is. The question is as to the proportion of true beliefs in the outputs of *this* habit over a sufficiently large spread of appropriately varied cases, in typical circumstances. In other terms, the question is as to the reliability of forming a belief *like this* on the basis of a perceptual presentation *like this*, where the relevant respects of likeness are determined by the constitution of the function realized in *this* mechanism.

As we have already seen, this point can be made more sharply with inferential belief-forming mechanisms, since there the function can be more easily specified, at least for deductive inferences. If I form the belief that *Jim will come to the party* on the input of the beliefs that *Jim will come if he is well* and that *he is well*, then the belief formed here will be reliably formed provided my principle of inference is conditionally reliable—such as to lead to truths from truths, and the input beliefs have been reliably formed. The first

of those two conditions will be realized if the principle that constitutes the function realized in the mechanism that gave rise to this belief is modus ponens. In that case the relevant type is such that not only is there a large proportion of true beliefs generated by processes of that type that take true beliefs as inputs; the process invariably yields true beliefs from true belief inputs.

Nondeductive inference, as we have seen, presents a much messier picture. Here, in order to be significantly reliable, a function must be sensitive not only to formal properties of the argument, e.g., that it is a generalization from instances, but also to a variety of more substantive considerations—the character of the sample if it is a case of generalization, various bits of relevant background knowledge, the field of competing explanations if it is an explanatory inference, and so on. Still, whatever the complexities of nondeductive-inference functions, the basic point is the same. What the epistemic status of the particular belief depends on, according to reliabilism, is the conditional reliability of the operative mechanism, i.e., the extent to which the function realized by that mechanism would yield true belief outputs from true inputs in a suitable spread of cases, plus the epistemic status of the input. Or, to put it in terms of processes, the crucial issue is the conditional reliability of the process of going from input to belief output in accordance with that function, along with the epistemic status of the inputs.

So when we think of the reliability of belief-forming processes in this way, we are no longer faced with an indefinitely large multiplicity of types among which we have to find some way of making a choice. With the illumination shed by this way of construing belief-forming processes, the "Problem of Generality," as construed by Feldman, Plantinga, and Pollock, dissipates like mist before the morning sun. To be sure, it is still true that a particular process token is an instance of an indefinite variety of process types, including countless silly ones like *processes that take place on Wednesday* and *processes that take place in the shower*. But now that we think of belief-forming processes as the functioning of a mechanism that embodies an input-output function, we can ignore all that. The function defines the epistemologically relevant type, and we can forget about the rest.

VI

But though the plurality of candidates for the position of relevant type has been greatly reduced, aren't there still serious alternatives among which this "functional" conception of belief-forming processes does not choose? Are there not many different ways of carving the psyche up into distinct mechanisms or habits of belief formation? Depending on which of these we pick,

we will end up with one or another assignment of a particular process to a general habit, habits that differ as to degree of reliability. Go back to my formation of the visual belief that there is a maple tree in front of me. Can't I think of the habit involved as one of (*a*) coming to believe that there is a maple tree in front of me on the basis of a visual presentation with such and such features, or (*b*) coming to believe that there is a tree in front of me on the basis of sensory experience, or (*c*) coming to believe that a certain plant is spatially related to me in a certain way, on the basis of sensory experience, or . . . ? That is, it looks as if I may think of the habit activated as possessing any one of widely varying degrees of generality. And it may well be that these habits differ in degree of reliability. Presumably the reliability of the habit of forming maple-tree-in-front-of-me beliefs on the basis of visual presentations with just these features is much higher than the reliability of the more general habit of forming tree-in-front-of-me beliefs on the basis of some sensory experience or other. And does this not mean that we have still failed to pick out a unique relevant type?

NO. At least, we are not still confronted with that problem if the assumptions I have been making are warranted. To properly respond to the above objection I need to distinguish the input and output sides. For the latter we can safely assume that the content can be determinately specified. We only have a case of belief formation if we have a belief, one with a particular propositional content. That gives us our starting point. So if we have a case in hand at all, we are not free to specify the output end of the function in various ways. If the belief the formation of which we are considering is a belief that there is a maple tree in front of me, then that ties down the output side, and there is no scope for choosing between different ways of characterizing it. No doubt, that belief belongs to various wider classes of beliefs—*that some plant is in front of me*, *that something is in front of me*, *that some plant is spatially related to me*, and so on. But none of these specifications give us the precise or full account of the content of the belief in question. Hence whatever mechanism produced this belief can't be one that embodies a function for going from a certain input to a belief with one of those more general contents, for that is not the belief that was actually produced in this instance. So we need not worry about a latitude in how the output side of the operative function is specified.

The input side and the function involved are a bit trickier, because they, especially the function, are not so open to view. Just what features of the input are picked up by the mechanism and just how the function "uses" them to determine features of the output (i.e., just what the function is) is not so obvious. At this point it becomes clear that the current objection forces me to become explicit about a basic assumption of my approach to the problem, viz., its **psychological realism**. I assume that there is always (almost always?) a single answer to the question, "Just what mechanism, embody-

ing just what function, was operative in the generation of this belief?" I assume that just one way of generalizing from this particular sensory-input–belief-output relationship reflects the actual psychological dynamics of the situation. When I look out the window and form the belief that there is a maple tree there, there are, in the abstract, many functions that would yield a belief with that content from a visual presentation of that sort. The transition might be based on the leaf shape, the overall shape of the tree, the color of the leaves, the character of the bark, the size, etc., or some combination of such features. But I am assuming that in that case only one of these possibilities is realized. The mechanism that was operative embodied one of these ways of taking certain features of the concrete input rather than others as the ones that yield a belief with that content. Again, when my beliefs that John will come to the party if well and that John is well yield the belief that John will come to the party, there are many abstract possibilities as to the principle of inference involved. The mere fact that the inference exhibits a modus ponens form does not guarantee that this is the principle that was psychologically operative. I could have been utilizing a function that yields that belief on the basis of any beliefs about John, or on the basis of any set of beliefs one of which is a conditional, or. . . . Nevertheless, according to my psychological realism, exactly one of those possibilities is realized in this case. And whichever one is realized, it is the reliability of that function (or of the correlated mechanism or process) that is crucial for the epistemic status of the belief.

Like any form of realism this one can be opposed. One can doubt or deny that the psyche really is determinate in this way; and one can hold that we are therefore free to make any of indefinitely many choices in picking a determinate function for purposes of assessing a given belief for degree of reliability. And here as elsewhere epistemological motives for antirealism are prominent. It may be claimed that we lack the access to the details of cognitive processes that would be required to determine in each case just exactly what function is operative. One of Plantinga's objections to reliabilism is along these lines.

> Indeed, if, as Goldman suggests, the relevant type must be specified in psychological or physiological terms, we won't be able to specify any such types at all; our knowledge is much too limited for that.[23]

I can't enter into a full-dress defense of psychological realism in this essay. The main point I am concerned to make at this point is that the viability of a reliabilist theory of justification or knowledge hangs on the viability of psychological realism. If there is not an objective fact of the matter as to what input-output function is utilized in a given belief formation, then reliabilists are helpless before the Problem of Generality, and they may as

well pack their bags and go home. This is why I said that the issues raised by the current objection are so crucial. They go straight to the heart of what it takes for reliabilism to be a real possibility for epistemology.

But short of a full-dress defense I will say this. First, we should not unduly inflate the epistemological requirements for psychological realism about belief-forming mechanisms. To be sure, if our position vis-à-vis such mechanisms were one of total ignorance, the game would not be worth the candle. But we should also be alive to the point that it is highly reasonable to suppose that there are many objective facts we will never know about, and even facts we are incapable of knowing about. Details of the past history of humanity, the earth, the solar system, and, more generally, the universe present many examples of this. Hence human epistemic incapacities vis-à-vis *X*'s are not necessarily fatal to realism about *X*'s.

My second point is that our cognitive access to belief-forming mechanisms is not as scanty as my antirealist opponent makes it out to be. Although we, obviously, can't peer into a psyche with some instrument and observe the little input-output functions doing their thing, we are in a similar situation with respect to many other matters about which we know something; and the approaches we use in those cases are available here too. When it is a question of what function was operative in my own case, I often have a "participant knowledge" of this. Although my knowledge of my own input-output mappings is far from ideally complete, and although I am not immune from error in such matters, I typically can, by reflecting on what is going on, gain some significant degree of insight into what it was that led me to form a certain belief. In the perceptual cases, even if I can't spell out in detail the "atomic" perceptual cues that enabled me to recognize what I was looking at as a maple tree, still I have some insight into the "look" of the object that tipped me off; and I can recognize relevantly similar looks when they recur. In inferential cases I can often, if I am analytically inclined, formulate the principle(s) of inference on which I was relying; or at least I can recognize relevantly similar inferences when they occur. As for third-person cases, if I want to know what function was operative in Sam's inference about John's coming to the party, I can put him to the test in other cases with and without a modus ponens structure, and with or without the other features mentioned above. In that way I can try to find consistent patterns in the way he draws conclusions; and although success is not guaranteed, I might reach some fairly solid results. The perceptual cases are again more difficult, but there too I can vary input-output relationships in such a way as to give myself a chance to find consistent, fairly stable patterns of relating perceptual-presentation features to belief content. These techniques yield less than maximally conclusive results for several reasons. For one thing, a person might be utilizing a given function on one occasion, even though she doesn't do so regularly,

often, or consistently. Habits of belief formation, like other habits, can be more or less stable, and they definitely are subject to change. For another thing, no matter how many competing hypotheses we have eliminated as to what the function is on a given occasion, there are always more looming on the horizon. Nevertheless, by using techniques such as these, we are considerably better off than blankly ignorant as to what function is operative in a given case of belief formation.

Another point about psychological realism. I don't want to overstress the determinacy and precision of belief-forming functions or, for that matter, belief contents. Indeterminacies to which all psychological states and operations are heir are to be found here as well. Some, or all, of my belief-forming habits may be such that there is some leeway as to exactly where certain perceptual features must be on relevant dimensions (of size, color, pitch, etc.) in order to generate a certain belief content. Even if the input sensitivity is perfectly precise, there may be some looseness in the way in which input features determine belief content. Exactly the same features might at one time yield the belief that a maple tree is present and at another time (with the "same mechanism" operative) yield the more specific belief that a sugar maple is present. Indistinguishable samples might on different occasions yield general beliefs with somewhat different statistical parameters or somewhat different degrees of confidence. And so on. The psychological realism I espouse is committed only to a degree of determinacy of belief-forming mechanisms that is sufficient to make it worthwhile invoking them in thinking about the reliability of belief formation as well as many other matters.

Here are two other complexities that would have to be recognized in an adequate theory of belief-forming processes and an adequate reliabilist epistemology. First, more than one habit might be involved in a particular belief generation. My belief that it is my wife's car that I see parked in our driveway might be generated both by a perceptual mechanism that takes account of features of my current visual experience and by an inferential mechanism that takes as input the belief that she told me when I left in the morning that she would be at home all day. I don't think this kind of overdetermination poses any special difficulties for cognitive psychology, but it does require the reliabilist epistemologist to make a decision as to which mechanism is such that its reliability is crucial for the justificatory status of the belief. As I have said, I am not, in this paper, setting out to develop a reliabilist epistemology. But I will say, in passing, that if each process would have been sufficient by itself to produce that belief, it would seem reasonable for the reliabilist to hold that the belief is justified provided either of the processes is sufficiently reliable.

The second complication is this. I have been talking as if every belief is generated by a single momentary input-output mapping. But, as we all know,

some beliefs are arrived at only after a more or less extended period of deliberation, search for evidence or reasons, weighing considerations pro and con, and so on. How are we to fit that sort of thing into the picture I have been developing? Here I believe that it is primarily the psychologist, rather than the epistemologist, who has additional work to do. In developing the psychology of belief formation, the cognitive psychologist has to decide how to represent the structure of these extended deliberative processes. For one thing, more than input-belief output mappings are involved. The searches for relevant evidence and weighing of pro and con considerations are processes of a different character. Perhaps the thing to say is that the belief-forming process occurs only at the end of the deliberation and that when it occurs it is of the simple, momentary sort of which I have been speaking. Or perhaps some other construal would be preferable. In any event, I am happy to leave this issue to the cognitive psychologist. So far as I can see, a reliabilist epistemology could work with whatever account seems best from the standpoint of psychological theory.

Interestingly enough, Goldman, who, as we have seen, proposed the *input-output function* conception of belief-forming processes, has failed to take advantage of this idea to solve the generality problem. Here is what he says about that problem in the very essay in which that functional construal was advanced.

> A critical problem concerning our analysis is the degree of generality of the process-types in question. Input-output relations can be specified very broadly or very narrowly, and the degree of generality will partly determine the degree of reliability. . . .
>
> It is clear that our ordinary thought about process-types slices them broadly, but I cannot at present give a precise explication of our intuitive principles. One plausible suggestion, though, is that the relevant processes are *content-neutral*. It might be argued, for example, that the process of *inferring p whenever the Pope asserts p* could pose problems for our theory.[24]

There is no hint that the identification of the function involved in a particular belief acquisition itself serves to define the relevant type of process. The discussion in Goldman's *Epistemology and Cognition* goes beyond the above remarks, but, aside from using the propensity conception of reliability to dissolve the Single-Case Problem, the additional suggestion amounts to the following.

> But how is it determined, in each specific case, which process type is critical?. . . Let me advance a conjecture about the selection of process types, without full confidence. The conjecture is: the critical type is the *narrowest* type that is *causally operative* in producing the belief token in question.[25]

But if my remarks in this essay are on target, there is only one "type," i.e., only one realized function, that is *causally operative* in a given case, apart from cases of overdetermination. The function, and the process type it defines, can, of course, be characterized in various ways, as we have seen, because it has various properties, intrinsic and relational, beyond the constitution of the function that specifies its epistemically relevant character. But when we are interested in knowing what kind of process to check for reliability in order to assess the belief epistemically, the realized function that was actually operative in the belief generation gives us a unique answer to that (if we know enough about the particular case to take advantage of this possibility).

Why haven't philosophers seen this, even philosophers like Goldman who realize that a belief-forming process is, essentially, the operation of a realized input-output function? Again, I suspect that the reluctance to be this psychologically realist plays a major role here. If one simply talks of "processes," each of which belongs to indefinitely many classes or types of processes, one is not making such strong psychological assumptions. But, as I said above, if one is not prepared to be that realist about the psyche, one should quit trying to be a reliabilist in epistemology.

VII

This completes my central task in this paper—exhibiting a psychological, nonepistemic basis for identifying a unique type in terms of which a particular case of belief formation can be assessed for whether the belief was reliably formed. It is not the case that all type assignments based on properties of the process are on a par. In every case there is one that reflects the actual dynamics of the belief formation, and, naturally, it is this one to which we should look if we are interested in how that particular belief formation stacks up with respect to reliability.

I have said more than once that it is no part of my intention here to enter into the issues that arise when one tries to build on this result to develop a reliabilist account of justification and/or knowledge. Nevertheless, the above results can be used to throw light on some further objections to that enterprise. In this section I will say a few words along those lines about what Feldman calls the "No-Distinction Problem."

You will remember that Feldman presents that problem as follows.

> A very broad account of relevant types of belief-forming processes leads to what we may call "The No-Distinction Problem." This arises when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type.

> For example, if the relevant type for every case of inferring were the type "inferring," then (RT) would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type.[26]

Of course the reliabilist will seek to make type assignments such that it is not implausible that all beliefs produced by tokens of a single type have the same epistemic status. The question for me is as to whether epistemically relevant types picked out in the way I have been suggesting pass the no-distinction test. Are actual belief-forming *functions* sufficiently homogeneous epistemically? I am not going to try to answer that question in this paper. Instead I will point out how the reasons given by Feldman and Plantinga for doubting that any types satisfy this requirement can be seen to fail if we restrict ourselves to psychologically realist types. And I will end with a few considerations that suggest that the belief-forming functions in the human psyche might well be epistemically homogeneous. In this discussion I will assume that the functions in question are maximally specific, in that any difference in input that is registered by the function indicates a different function.

One case used by Feldman, following Goldman, has to do with visually recognizing a mountain goat. Feldman has been making the obvious point that so broad a type as *visual belief formation* fails the no-distinction test. He then cites Goldman's statement that seeing a nearby object is a different process from seeing a distant object. Feldman supposes that this difference has to be drawn in terms of external conditions of observation.[27] But that is not so. The relevant features of the sensory experience will obviously be different in the two cases, and so a different function will be activated. At a later stage he considers borderline cases such as often confront an umpire calling balls and strikes. "Some of the objects the person sees may clearly have the property in question while others do not. As a result, some beliefs to the effect that the object has the property may be better justified than others."[28] And so the products of "the process" are epistemically heterogeneous. But, again, it is plausible to suppose that the subject is sensitive to differences between cases that are visually clear and cases that are not. Hence there will be at least two different functions involved.

I find that all of Feldman's cases can be plausibly disposed of in this way, once we think of relevant types in terms of the input-output functions. Of course, this doesn't show that all belief-forming functions are epistemically homogeneous, any more than Feldman has shown that many of them are not.[29] As previewed, I will end by presenting some considerations that are relevant to the consideration of this question.

First it is clear that there is no possibility of epistemic heterogeneity for

deductively valid inferential functions, like modus ponens, that are logically guaranteed to yield a true conclusion from true premises whatever the subject matter. Outside that realm we don't have the same sort of guarantee. But it still seems plausible that the kinds of functions we have been discussing don't suffer from epistemic heterogeneity. As for nondeductive inference, if we were to try to treat this formally on the model of deduction, we would run into the No-Distinction Problem. Think of the formal inductive function—producing the belief that *Most F's are G's* on the input of one hundred *F*'s that are *G*'s without any *F*'s that are not *G*'s. Obviously this works better with some *F*'s and *G*'s than with others. It gives us a higher reliability if *F* is *golden retriever* and *G* is *affectionate* than it does if *F* is *book* and *G* is *has a red cover*. But it seems clear that the nondeductive inference functions generally internalized are not of this sort. They are, rather, designed to be sensitive to peculiarities of the subject matter as well as to the kind of inference involved. Again, an "argument-to-the-best-explanation" function will be sensitive to a good deal of specific information about the explanandum as well as about the competing explanations and the factors that affect their comparative assessment. And as for perceptual belief functions, so long as they restrict inputs to one set of experiential features, perhaps together with relevant background beliefs, and involve a single way of mapping these onto belief output content, there would seem to be no way in which some subset of their uses would involve a significantly higher proportion of true beliefs than others.

But since, as I have already been at pains to point out, our knowledge of the details of human belief-forming functions is imperfect at best, we cannot be sure that there are no functions that founder on the No-Distinction Problem. Perhaps some, or even many, people instantiate purely formal nondeductive inference functions that are defective in this way. Moreover, as the editor of this issue pointed out to me, even if no human belief-forming functions are epistemically heterogeneous, there are surely possible subjects that instantiate such functions. And doesn't a reliabilist epistemology aspire to handle possible as well as actual cases? Leaving that last issue aside, I will conclude by discussing the possibility of a single humanly realized perceptual function that fails the no-distinction test by virtue of delivering a belief with a certain propositional content on any one of a variety of inputs.

Think of the familiar point that there are various ways which the members of a kind look. Consider 'dog', 'tree', 'pine tree', or 'house'. There is an enormous diversity in the kinds of perceptual presentation that a particular person will regularly take as indicating that a perceived object is a dog or a tree or a house. Not all dogs or all houses look just alike, not by a long shot. And it may well be that my identification of something as a dog by the use of one of these component inputs is much more reliable than others.

Some doggy looks are more similar to wolf looks or fox looks than other doggy looks. Some pine trees are harder to distinguish from other conifers than other pine trees. And if it is correct in cases like this to count all my ways of going from the look of *X* to the belief that *X* is a dog, or a pine tree, as exercises of a single habit, then if we determine the reliability status of a particular belief by determining the reliability of the habit the activation of which gave rise to that belief, it looks as if we are stuck with a not inconsiderable number of cases in which the epistemic assessment of the particular belief will be indeterminate because of the heterogeneity of the habit that produced that belief.

But this is a serious worry only if it is a serious possibility that there are unitary belief-forming habits of this sort. And that is dubious. It is by no means obvious, from introspection and other ordinary means of access, just what to say about such cases. Where there is a single unified functional relationship that we can formulate and that is plausible to regard as the content of an actual belief-forming habit, as with perceptual color attributions, then we have an obvious reason for attributing a belief-forming habit that embodies that rather broad function to people. But here we have no such reasons available. Why should we suppose that anyone possesses a *single* function that includes all her ways of perceptually identifying a dog or a pine tree. If we have any basis for doing so, it will be a more recherché theoretical basis. And theoretical considerations would seem to tend in the other direction. After all, the only reason for supposing there to be a unified function here is the commonality in the content of the belief output. Aside from the fact that the inputs include features of perceptual presentations, that is the only thing that stretches over the whole territory. The particular experiential features vary widely, and the ways in which these features are mapped onto belief content vary correspondingly. Hence we have only the slenderest of excuses for supposing there to be a single function. Given that the basic notion of a belief-forming function is tied to *a certain way of mapping input features onto belief output content*, and given that many ways are lumped together here in the alleged single complex function, there are strong reasons for positing a large number of perceptual dog-recognition functions, and, in the absence of more compelling reasons on the other side, that is the best choice, apart from the exigencies of epistemological theory.

Hence I am not inclined to worry about the "No-Distinction Problem" in cases like this. To be sure, for any perceptual belief-forming function, no matter how specific, there will be some variation in experiential inputs. Even when I am looking at the maple tree in my front yard from exactly the same distance and angle and at the same time of day, there will still be minor variations from the quality of the light, the condition of the tree, and so on. But we must remember that the effective inputs to a habit, what it "takes account

of" in forming the belief, are *abstract* features of the presentation, not that presentation in all its concreteness. Hence the inputs can be the same through considerable variation in the total character of the presentation.

VIII

To sum up. I have suggested that we think of belief formation in a psychologically realistic way, as involving an input-output mechanism (habit) that yields belief outputs as a certain function of relevant features of inputs. If we do so, we can escape from the dilemma of what choice to make, for epistemic assessment as to reliability, of the type to which a particular belief-forming process belongs. That choice is settled for us by the identity of the function involved in the belief formation in question, for that function is something that possesses a built-in generality. In other words, a particular belief formation is the activation of a general mechanism (habit) that operates in accordance with a certain function. The mechanism (habit) *is* the psychological realization of that function. On this basis we can say what it is for a belief to be reliably formed.

> I. **A belief is reliably formed if and only if it was formed by the activation of a reliable belief-forming habit.**

And:

> II. **A belief-forming habit is reliable if and only if it would yield a high proportion of true beliefs in a sufficiently large and varied run of exercises in situations of the sorts we typically encounter.**

In conclusion, let me remind the reader of what I have and have not set out to do in this essay. My sole concern has been to explicate concepts of reliability that are usable in reliabilist theories of epistemic justification and of knowledge. I have not presented particular theories of this sort, much less defended or attacked them. My aim has been only to remove what are widely believed to be internal conceptual difficulties in such theories. I have been moved to engage in this enterprise by the fact that those who have criticized reliabilist epistemologies on the grounds of conceptual incoherence, and those who have responded to such criticisms as well, have not been alive to the way of understanding reliability that I have been developing here. My suggestion is that if we think of reliably formed belief in the way I have laid out in this essay, we will be able to dispose of the supposedly fatal difficulties in the concept of reliability put forward by criticisms like those of Feldman. That would leave the way clear for an evaluation of the merits and demerits of reliabilist epistemologies.

1. See Alvin Goldman, "What Is Justified Belief?" in *Justification and Knowledge*, ed. George Pappas (Dordrecht: D. Reidel, 1979), and *Epistemology and Cognition* (Cambridge, Mass.: Harvard University Press, 1986).
2. See Marshall Swain, *Reasons and Knowledge* (Ithaca, N.Y.: Cornell University Press, 1981).
3. See Frederick Schmitt, *Knowledge and Belief* (London: Routledge, 1992).
4. Is this a *conceptual* claim or not? Is it the view that the *concept* of justified belief is the *concept* of a reliably formed belief? Or is the concept of justified belief explained in some other way, reliable belief formation being put forward as a sufficient (and possibly necessary) condition for justified belief as so construed? It could be either. Goldman is quite explicit that reliability functions in the second way for him, whereas in Swain we have what looks like the claim that what we *mean* by saying that a belief is justified is that it is reliably formed (to greatly oversimplify Swain's actual formulation).
5. Goldman, "What Is Justified Belief?" 13.
6. The sequel will illustrate this, at least with respect to Goldman.
7. See Richard Feldman, "Reliability and Justification," *The Monist* 68 (1985): 159–174.
8. For a valuable discussion of the issues involved here, see Goldman, *Epistemology and Cognition*, sec. 5.5.
9. Feldman, op. cit., 159–160.
10. Alvin Plantinga, "Positive Epistemic Status and Proper Function," *Philosophical Perspectives* 2 (1988): 28–29. See also John Pollock, *Contemporary Theories of Knowledge* (Totowa, N.J.: Rowman and Littlefield, 1987), 116–121.
11. Feldman, op. cit., 160–161. Plantinga (op. cit., 29–30) echoes this presentation of Feldman's, and Pollock (op. cit.) makes similar claims.
12. Plantinga and Pollock give a similar assessment.
13. See Feldman, op. cit., 168.
14. Goldman, "What Is Justified Belief?" 11.
15. Ibid., 13–14. The reader will note that these two formulations differ with respect to whether it is *truth* or *justification* that is thought of as transferred across the inference. That does not mean that Goldman has contradicted himself. It can be consistently held both that to be "conditionally reliable" a dependent belief-forming process must yield *true* outputs from *true* inputs *and* that this requirement must be met in order for a dependent process to yield a *justified* belief from *justified* inputs. And yet it would seem unwarranted to require this for the justification of an output. That would imply that nondeductive inference never yields justified conclusions from justified premises; for such inference does not always yield true conclusions from true premises. (To be sure, Goldman presents conditional reliability of the process only as a *sufficient* condition for getting justified conclusions from justified premises. But since he mentions no other sufficient condition, one is naturally led to suppose that he takes the condition to be necessary as well.) However, these problems can be ignored for my purposes in this paper.
16. A skeptic might question whether we know that such situations are "clearly atypical." How do we know, she might ask, that we are not always in a Cartesian demon situation? It is no part of my intention here to answer, or otherwise discuss, skepticism. I will only point out that this last qualification to my account, along with all the rest of my suggestions, presupposes the falsity of radical skepticism. My proposals presuppose that we do have some considerable ability to determine when the conditions specified in those proposals are and are not satisfied. It is worth noting that if radical skepticism were correct, everything else would collapse, along with my account of reliability. Skepticism isn't a problem for it in particular.
17. See David M. Armstrong, *Belief, Truth, and Knowledge* (Cambridge: Cambridge University Press, 1973), 170.

18. In the course of his discussion Armstrong does put certain restrictions on H, but none of them have the effect of limiting the extent of the process.
19. Goldman, "What Is Justified Belief?" 12–13.
20. Ibid., 11.
21. I am speaking here of a *function* in the mathematical sense, the sense in which, to take a very simple example, addition is a function. Given any two or more numbers, the addition function will yield a unique output as their sum. A function in this sense is, of course, something abstract. That is why I had to specify that a belief-forming process involves the *activation* of a *psychologically realized* function, not just the function as a denizen of logical space.
22. The details of this account will vary with one's favored theory of perception. Since I don't want to get into those issues here, I am striving for maximum neutrality. Nevertheless, my talk of "presentations" reflects my attachment to a theory of appearing, according to which perceptual experience consists most basically of objects, usually external physical objects, *appearing* to one in certain ways. As a result, sense datum theorists, adverbial theorists, and conceptual-propositional theorists of perceptual experience will not like my way of putting it. I believe, however, that the points I am making here concerning the features of belief formation that are relevant to assessments of reliability are neutral with respect to different accounts of perceptual consciousness. I would invite those who take exception to my formulations to restate what I am saying in their favorite terms.
23. Alvin Plantinga, *Warrant: The Current Debate* (New York: Oxford University Press, 1993), 199.
24. Goldman, "What Is Justified Belief?" 12.
25. Goldman, *Epistemology and Cognition*, 50.
26. Plantinga and Pollock add an extra twist to this formulation. They both hold that the reliabilist needs to find a type that is homogeneous in the sense that there is no subtype of that type the reliability of which differs from the larger type. (See Plantinga, "Positive Epistemic Status and Proper Function," 30, and Pollock, op. cit., 119.)
27. See Feldman, op. cit., 163.
28. Ibid., 164–65.
29. He doesn't claim to have done so.