BELIEFS AND SUBDOXASTIC STATES*

STEPHEN P. STICH†

University of Maryland

It is argued that the intuitively sanctioned distinction between beliefs and non-belief states that play a role in the proximate causal history of beliefs is a distinction worth preserving in cognitive psychology. The intuitive distinction is argued to rest on a pair of features exhibited by beliefs but not by subdoxastic states. These are access to consciousness and inferential integration. Harman's view, which denies the distinction between beliefs and subdoxastic states, is discussed and criticized.

This is a paper about a distinction, one which is deeply imbedded in our everyday, pre-theoretic thinking about human psychology. It is a distinction which separates beliefs from a heterogeneous collection of psychological states that play a role in the proximate causal history of beliefs, though they are not beliefs themselves. I will call the states in this latter collection *subdoxastic* states. The distinction is an intuitive one, in the sense that when confronted with descriptions of various imaginable states it is intuitively clear that some are to be counted as beliefs while others are not. Part of my project in this paper is to explain these intuitive decisions by noting some of the characteristics that we ordinarily take beliefs to have which are lacking in subdoxastic states. Since these intuitive judgements mark part of the boundary of our pre-theoretic concept of belief, an explanation of the principles underlying the intuitions may also be viewed as part of an analysis of our ordinary concept of belief. It is this analytic project that occupies the early sections of the paper.

The intuitive distinction between beliefs and subdoxastic states is of interest quite apart from the insight it promises into our ordinary notion of belief. For much of the best recent work in cognitive psychology has ignored it. The implication, albeit a tacit one, of this neglect is that the distinction entrenched in intuition does not mark a psychologically interesting boundary. Now there is surely no *a priori* argument for rejecting this view. There is no reason to expect that every distinction embedded in our intuitive, pre-scientific psychological

*Received January, 1978.

†I am indebted to Clayton Lewis and Robert Cummins for helpful criticism of the views set forth in this paper.

Philosophy of Science, 45 (1978) pp. 499–518.

concepts reflects a distinction worth preserving in an empirically motivated psychological theory. What is unsettling about the neglect of the distinction in recent psychology is that it seems to be largely undefended. The fact that pre-theoretically we recognize the distinction as readily and consistently as we do is surely evidence that there is some basis to the distinction, some property or cluster of properties that we tacitly take account of when we classify a given example as a belief or as a subdoxastic state. The properties which we exploit in drawing the distinction may be psychologically uninteresting. But to be confident of this one would like to know what these properties are and why they can be safely ignored in psychological theorizing. The psychologists who ignore the distinction do not inspire confidence in this way. Such limited philosophical discussion as there has been of the topic is not reassuring. Gilbert Harman, in his recent book [7], sketches an argument that might be construed as an attack on the psychological significance of the distinction. So construed, the argument is singularly unconvincing. The latter sections of my paper will be devoted to a critique of Harman's view. I will also try to defend the view that the boundary between beliefs and subdoxastic states marks an important psychological distinction which serious cognitive psychology neglects to its peril.

1. Our first job is to focus in on the particular intuitive distinction I have in mind. Let me begin a bit obliquely by noting the shadow the distinction casts in traditional epistemology. It is obvious that many of our beliefs are inferred from other beliefs. But, as epistemologists were quick to note, there must be some beliefs which are not inferred from others, on pain of circularity or infinite regress. Since Descartes, the orthodox position has been to take as non-inferential those beliefs that concern the believer's own cognitive states—beliefs about what she believes, what she seems to be perceiving, what she seems to remember, etc. In recent years, the orthodoxy has been challenged by philosophers who advocate a much more expansive account of non-inferential belief. The expansionists urge that most perceptually-based beliefs about our immediate surroundings are non-inferential. Thus, while agreeing that our beliefs about our current cognitive states are non-inferential, the expansionists would also include among the non-inferential beliefs the belief that there is a pig nearby, if acquired under unexceptional circumstances as the result of a pig coming into view a few meters away. (cf., for example, [2], pp. 115 ff. and [4], p. 159.) We need not pause here to trace the moves in the debate between the expansionist and the advocate of orthodoxy. For our purposes it is important only to note that,

despite their differences, both sides in the dispute recognize that some beliefs must be non-inferential. These non-inferential beliefs may serve as the premises for inferences but not as the conclusion. The class of non-inferential beliefs marks the boundary between two different sorts of psychological states. On one side of the boundary are beliefs, both inferred and non-inferential. On the other side are the psychological states which, though not beliefs, are part of the causal process leading to belief formation. Typically, epistemologists have been silent on the nature and workings of the mechanisms which underlie belief formation. Their silence, no doubt, was rooted in the sensible suspicion that the study of subdoxastic mechanisms is more properly the province of psychology.

2. To sharpen our intuitive feel for the distinction between beliefs and subdoxastic states, let us look at a few examples of the sort of states that fall in the latter category. Consider first the offering of grammatical intuitions. Given a few well-chosen examples of grammatical sentences and a few well-chosen examples of ungrammatical ones, speakers are able to go on to classify new examples as either grammatical or ungrammatical with impressive consistency. This ability seems to be largely independent of previous formal grammatical training. Naive subjects (particularly bright ones) quickly get the point. On the other hand, the ability plainly does depend on some elaborate and little-understood system of psychological states and processes which is gradually built up in the process of language acquisition. Moreover, it seems plausible to speculate that the psychological mechanisms which underlie the offering of grammatical intuitions also play a central role in the much more important business of speech production and comprehension. (For a general discussion of grammatical intuitions see my [14] and [15]; for some critical reaction, see [3] and [5].)

Now suppose that a grammatically naive subject has just been taught by example to apply the labels "grammatical" and "ungrammatical." We give the subject a new sentence and ask whether or not it is grammatical. The subject replies—and presumably believes—that it is grammatical. But if we ask him how he made the judgement and came to hold that belief, he will be at a loss to say. On hearing the sentence, he simply comes to believe that it is grammatical. The belief is non-inferential. There is, of course, a mechanism of some complexity mediating between the subject's hearing the sentence and the formation of the belief that it is grammatical. And, while we know little in detail about the workings of this mechanism, it is plausible to speculate that the mechanism exploits a system of psychological

states which serve to store information about the grammar of the subject's language.[1] If this speculation proves accurate, then these states which store grammatical information are a prime example of subdoxastic states. They play a role in the proximate causal history of beliefs, though they are not beliefs themselves. I have, of course, given no argument that these states are not beliefs. My claim, rather, is that this is intuitively obvious. After we have seen a few more examples of subdoxastic states, I will try to dissect out those aspects of our concept of belief which are responsible for the intuitions.

3. For a second example of subdoxastic states, let us attend to the complicated matter of depth perception. Our ability to judge the relative distances of objects in our visual field rests on a bewilderingly complex set of factors. Perspective, size, surface texture, the perception of edges and corners, occlusion, illumination gradients and stereopsis all play a role. (For a survey of this topic, see [6], chapter 13.) Some of the information available to our depth perception mechanisms would appear to be redundant. We can, for example, make depth judgements with one eye closed, or about the objects in a painting. In both cases stereopsis is irrelevant. Julesz, in an interesting series of experiments, has shown that binocular disparity alone is enough to bring about perception of depth. [10]. The experiment employed random patterns of black and white dots in a 100×100 cell matrix. The randomness of the patterns served to guarantee that no depth information was conveyed by the patterns themselves. When two identical copies of a random dot pattern are presented to a subject, one to each eye, no depth is perceived. However, if a section of one copy of the pattern is displaced laterally, there is a clear perception of depth, with the moved section appearing either in front of or behind the rest of the pattern, depending on whether it was moved toward the nose or toward the ears. When a subject in Julesz's experiment is presented with a pair of dot patterns one of which has a laterally displaced section, he will report—and presumably believe—that some of the dots seem to be in front of or behind

[1]Speakers might, for example, have a system of states which serves to store or internally represent each of the rules of the grammar. In judging whether a sentence is grammatical, the relevant psychological mechanism might try to produce a derivation of the sentence at hand, aided by some efficient derivation finding heuristic. If it succeeds, it sets in motion a process which ordinarily leads to an affirmative answer. I think this sort of "analysis by synthesis" strategy, though surely logically possible, is wildly implausible as an empirical hypothesis. It does, however, have the advantage of being a particularly straightforward example of the way internally represented grammatical information might be exploited. For a much more subtle illustration of how grammatical information might be represented and utilized, cf. [16].

the rest. Asked why he believes this, he will be unable to say. Some dots just do look to be in front of the others, and the belief that they do is non-inferential. Underlying the belief there is, no doubt, a complex psychological mechanism which serves to measure the degree of binocular disparity and to use this information in the production of an appropriate belief about apparent relative depth. If, as seems inevitable, the process leading from retinal stimulation to belief involves various psychological states which represent features of the retinal images, then these states are subdoxastic. Though they are causal antecedents of belief, there is a strong intuitive inclination to insist that they are not beliefs themselves.

4. Let us look now at a final example of subdoxastic states. In a series of experiments, E. H. Hess ([8]; see also [9]) presented male subjects with a pair of almost identical photos of a girl. One of the photos had been retouched by enlarging the size of the girl's pupils. The subjects regularly reported (and presumably believed) that the girl in the retouched photo appeared more attractive. However, subjects were quite unable to say why the girl in the retouched photo was more attractive; indeed, they were generally unable to identify any specific differences between the original and retouched photos. Their belief that the girl is more attractive in one photo than in the other is non-inferential. Plainly, there is some cognitive mechanism which detects the enlarged pupil size and which gives rise to the belief that one photo is more attractive than the other. As in the previous examples, we know little about the detailed functioning of this mechanism. However, it is plausible to suppose that there are psychological states which serve to record information about pupil size, and that these states play a role in the process that leads to belief formation. If so, these states are another example of subdoxastic states. They play a role in the proximate causal history of beliefs, but there is a strong intuitive inclination to deny that they are beliefs themselves.

5. Having surveyed some examples of states that fall on the subdoxastic side of the belief-subdoxastic state distinction, let us now consider what the basis may be for our strong intuitive inclination to insist that the sorts of states we have been describing are not beliefs. I think there are two rather different properties that we ordinarily take to be characteristic of belief and which are lacking in the examples of subdoxastic states. The first of these is actually a cluster of properties that revolves around the sort of *access* we ordinarily have to the contents of our beliefs; the second might be characterized as the *inferential integration* of beliefs. I will consider them in turn.

It is ordinarily the case, for typical or paradigmatic examples of belief, that adult subjects can report the contents of their belief. Thus, if a subject is psychologically (and physiologically) normal, inclined to be cooperative and has no motivation to deceive us, then if she believes that *p* and is asked whether *p* is the case, she will generally say that it is. (The subject must, of course, be asked in a language she understands and she must be paying attention. Hereafter I will take this to be built into the notion of asking.) Normal cooperative adult subjects can also tell us whether they *believe* that *p*, if we ask them. The two abilities are distinct. Toddlers, for example, are able to assent to *p* if they believe it, but are often not able to answer the question: "Do you believe that?" Nonetheless, the two abilities run in tandem in normal adults, and it is hard to believe that they are not intimately connected.

Both of these abilities are themselves associated with an ability to become aware of or to be conscious of the contents of one's beliefs. Suppose, for example, that a subject believes that *p* but fails to satisfy one or another of the additional conditions which, conjoined with belief, generally guarantee that a subject will assent to *p* if asked. The subject may be temporarily paralyzed and thus unable to assent to anything. Or he may have a strong desire to mislead his questioner, or simply wish to say nothing. Still, under these circumstances, if we ask a subject whether *p* is the case, he will generally have a certain sort of characteristic experience which, as best I can discover, has no standard description in English. Some philosophers have labelled the experience "having the occurrent belief that *p*." One might also describe the experience as being aware that *p* or being conscious that *p*. Often, but perhaps not always, the experience of having an occurent belief is accompanied by (consists in?) a perceived inclination to respond to the inquirer's question, or perhaps a response which is thought but not spoken. I do not propose to attempt any thoroughgoing phenomenological account of the experience of having an occurrent belief. Nor, in fact, am I much concerned whether it is a single sort of experience or a group of related ones. My point is the (I hope) uncontroversial one that in typical cases of belief a subject will have a certain sort of characteristic conscious experience when his attention is suitably directed to the content of the belief. While little is known about the psychological mechanisms responsible for the experience of having an occurrent belief, it is certainly plausible to speculate that the processes underlying this conscious experience and the processes underlying a subject's assent to a proposition he believes (under normal circumstances) are in important ways interconnected.

I have been claiming that access to the contents of our beliefs is a general characteristic of beliefs. However, there is a *prima facie* exception to this generalization viz. the unconscious beliefs that figure prominently in psychoanalytic theory. While I admit to some qualms about how seriously psychoanalytic theory ought to be taken, I do not think I need take on all of psychoanalytic theory to defend the claims I want to make about a subject's access to the contents of her beliefs. For all of my claims can be construed as subjunctive conditionals about what would happen (or would be likely to happen) if a subject were normal and otherwise suitably situated. And psychoanalytic theory can be viewed as postulating a psychological mechanism capable of interfering with the ordinary process leading from belief to assent or to conscious awareness. So to protect our generalizations about access, we will have to add to the conditions in the antecedent of our subjunctive conditional a clause specifying that there are no psychological mechanisms at work blocking the ordinary process leading from belief to assent or conscious awareness.

Now one of the ways in which our first two examples of subdoxastic states differ from beliefs is just that subjects have no access to them. People do not assent to a statement of a rule of their grammar as they do to a statement of the contents of a belief. Nor do they have any conscious awareness of the rules of their grammar. Indeed, if they did writing a grammar for a language would be a far less arduous business. Similarly, we cannot report, nor are we consciously aware of any information about binocular disparity. Awareness stops at the level of apparent comparative depth. Of course, it might be the case that these seemingly subdoxastic states are actually analogous to subconscious beliefs in that subjects would have access to them but for the intervention of some presently unsuspected psychological mechanism which blocks the ordinary processes that facilitate access. This possibility reflects something important about our pre-theoretic concept of belief. We might be willing to classify an apparent subdoxistic state as a belief if there were evidence that a subject's access is actively blocked by some mechanism. But we would, I think, be much more reluctant to countenance a special category of beliefs which are by nature not open to conscious awareness or reporting. It is quite central to our concept of belief that subjects under ordinary circumstances have access to their beliefs.

Our third example of a subdoxastic state poses some special problems. In Hess' experiments, subjects were not consciously aware that the pupils in one photo were larger than those in the other. However, at least some of the subjects might have become aware of it if their attention were suitably directed by, say, asking them.

In this respect, the state which serves to represent the information that one pupil is larger than another is analogous to quite unexceptional cases of belief. For we are ordinarily quite unaware of most of our beliefs, and the experience of having the belief occurrently is provoked when our attention is directed to the content of the belief. I think this similarity between unproblematic cases of belief and the state that we have hypothesized in the Hess example is reflected in our intuition. For of our three examples, our intuitive inclination to rule that a state is not a belief is weakest in this case. But still, I think there is a strong inclination to resist calling the supposed state a belief. An attractive alternative is to say that, though not a belief itself, the state can give rise to a belief under certain circumstances. And thus the role of drawing a subject's attention to the relative size of the pupils is *not* analogous to what goes on when, for example, we ask a subject (who has been thinking about other things) what her mother's maiden name is. In the latter case, the question serves to make occurrent or bring to consciousness a belief which the subject has had all along. In the former case, however, questioning the subject about relative pupil size serves rather to instigate a process of belief formation in which, perhaps, the pre-existing subdoxastic state plays a role. If I am right that this alternative account of what is going on sits better with our intuition than the account which assimilates the Hess case to standard examples of belief, then we are left with the question of why this should be so. What is it about the state in the Hess example that makes us reluctant to treat it as a belief? I think the answer is that this state shares with our other two examples of subdoxastic states a sort of inferential isolation from the body of our accessible beliefs. This is the topic to which I now will turn.

6. It is characteristic of beliefs that they generate further beliefs via inference. What is more, beliefs are inferentially promiscuous. Provided with a suitable set of supplementary beliefs, almost any belief can play a role in the inference to any other. Thus, for example, if a subject believes that p and comes to believe that if p then q, he may well come to believe that q—and do so as the result of an inferential process. In addition to the well-integrated network of potential deductive inferences, beliefs also generate other beliefs via inductive inference. So there is generally a huge number of inferential paths via which a given belief can lead to most any other. It is in this sense that a person's beliefs are *inferentially integrated.*

Of course, the patterns of valid inference specified in deductive logic and those that would be specified by a theory of inductive logic, if there were such a theory, correspond in no simple way to

the inference pattern exhibited among a person's beliefs. We do not, for example, draw all logically possible inferences from our beliefs. Nor are all our inferences logically valid. What is more, there is no reason to suppose a priori that each logically permissible inference is in fact psychologically possible. And it is at least possible there are some beliefs that cannot be acquired by inference at all. The non-inferential beliefs of traditional epistemology would be an example, if indeed there are any such beliefs. However, none of these caveats detracts from our principal point about the inferential integration of beliefs: a person's body of beliefs forms an elaborate and interconnected network with a vast number of potential inference patterns leading from every belief to almost any other.

Now it is my contention that part of the reason we are intuitively inclined to say subdoxastic states are not beliefs is that subdoxastic states, as contrasted with beliefs, *are largely inferentially isolated from the large body of inferentially integrated beliefs to which a subject has access.* This is not to say that subdoxastic states do not play any role in inference to and from accessible beliefs, but merely that they are inferentially impoverished, with a comparatively limited range of potential inferential patterns via which they can give rise to beliefs, and a comparatively limited range of potential inferential patterns via which beliefs can give rise to them. This last remark may strike some as paradoxical. For, it might be protested, inference is a relation *among* beliefs. It is (one of) the ways beliefs generate other beliefs. And if subdoxastic states are not beliefs, then they cannot be *inferentially* related to anything. The objection is an important one, and its central claim—that inference is a relation exclusively among beliefs—is one that will play a prominent role in the following section. For the moment, however, let us grant that at least some of the routes by which subdoxastic states give rise to beliefs or other subdoxastic states can sensibly be taken to be instances of inference. Granting this assumption, the picture I am urging is this. Consciously accessible beliefs are embedded in an elaborate network of potential inferential connections with each other. Each belief is a potential premise in inferences to a vast array of further beliefs. By contrast, the inferential contact between subdoxastic states and beliefs is specialized and limited. When a subdoxastic state can serve as a premise in an inference to beliefs, there is only a narrow range of beliefs to which it may potentially lead. Similarly, when a subdoxastic state can result from an inference with beliefs among the premises, the range of beliefs that can serve in this capacity is restricted and specialized. If we think in terms of a cognitive simulation model, the view I am urging is that beliefs form a consciously accessible,

inferentially integrated cognitive subsystem. Subdoxastic states occur in a variety of separate, special purpose cognitive subsystems. And even when the subdoxastic states within a specialized subsystem generate one another via a process of inference, their inferential interactions with the integrated body of accessible beliefs is severely limited. Similarly, in all likelihood, the potential inferential connections among subdoxastic states in different specialized subsystems are extremely limited or non-existent. To get a clearer view of my thesis about the inferential integration of beliefs, let us see how it applies to our examples.

Even on the most generous assumptions about the inferential potential of the states which store grammatical information, these states are largely isolated from the body of a subject's beliefs. These grammar storing states can plausibly be assumed to play a role in the formation of beliefs about what has been said to a person (that is, in the process leading from auditory stimulation to comprehension). They can also, as we have seen, be presumed to play a role in the process leading to formation of beliefs about grammatical properties and relations—the sorts of beliefs expressed when an informant is offering grammatical intuitions. And let us suppose that the processes involved are properly viewed as inferential. Then the states storing grammatical information do have some limited inferential links with beliefs. But let us compare the state storing a grammatical rule with the explicit belief in that rule, say on the part of the linguist.[2] The linguist's explicit belief can enter into an almost endless number of inferences in which the competent speaker's subdoxastic state cannot participate. For example, if a linguist believes a certain generalization to the effect that no transformational rule exhibits a certain characteristic, and if he comes to believe a given transformation which violates the generalization, he may well infer that the generalization is false. But merely having the rule stored (in the way that we are assuming all speakers of the language do) does not enable the linguist to draw the inference. As another example, suppose that, for some putative rule *r*, you have come to believe that if *r* then Chomsky is seriously mistaken. Suppose further that, as it happens, *r* is in fact among the rules stored by your language processing mechanism. That belief along with the subdoxastic state will not lead to the belief

[2]There are a pair of assumptions being made here. First, I assume that the grammatical information stored by a competent speaker consists (in part) of rules of the speaker's grammar. As noted in fn. 1, I suspect this is a counterfactual assumption. Second, I assume that grammatical rules can be formulated in such a way that it makes sense to talk of believing a rule. More specifically, I am assuming that appending a formulation of a rule to '*S* believes that _____' produces a well-formed sentence.

that Chomsky is seriously mistaken. By contrast, if you believe (perhaps even mistakenly) that *r*, then the belief that Chomsky is seriously mistaken is likely to be inferred. It would be easy enough to marshall many more illustrations of the fact that the subdoxistic states which store grammatical information are largely inferentially isolated from beliefs.

The situation is, if anything, clearer for the states which process and store information about binocular disparity. Let us suppose that at a given stage in the process leading to the formation of judgements about apparent comparative depth in Julesz's experiments, a certain subdoxastic state represents the information that a dot on a certain part of the left retina is displaced five seconds further toward the nose then the similarly situated dot on the right retina. It may well be that a fairly complicated process of computation and inference is required for the formation of this subdoxastic state. The process might, for example, utilize and compare information about the distribution of dots on each separate retina in order to locate "corresponding" dots. It may also be the case that the process leading from our hypothesized subdoxastic state to the belief that the displaced dot appears in front of the others is itself inferential. On this assumption, the subdoxastic state does have a potential inferential path to an accessible belief. However, it is a vastly more restrictive path than beliefs have to each other. Contrast, for example, the *subdoxastic state* representing the information that a dot on a certain part of the left retina is displaced five seconds further toward the nose than the similarly situated dot on the right retina, with the *belief* that the dot in that part of the retina is displaced five seconds further toward the nose than the similarly situated dot on the right retina. The subdoxastic state can lead directly only to a restricted class of beliefs about apparent relative depth (and perhaps some other aspects of the visual field). By contrast, the belief, if supplemented by suitable additional beliefs, can lead to just about any belief. There is also a striking contrast in the ways other beliefs can *lead to* either the subdoxastic state or the belief. A subject might inferentially acquire the belief (that a dot on a certain part of the left retina . . . etc.) in numerous and diverse ways. He may be told that the dot is thus displaced by a person he takes to be trustworthy, and infer that it is on the basis of his belief about what his informant believes. Or he may infer it from beliefs formed by observing the readings on certain test instruments. Indeed, most any other belief, say the belief that *p*, can inferentially give rise to the belief about the relative positions of the retinal dots, provided that the subject also believes a proposition of the form *if p then d*, where *d* is the proposition

that a dot on a certain part of the left retina . . . etc. On the other hand, it is most likely the case that there are *no beliefs at all* which can lead inferentially to the subdoxastic state that represents the fact that *d*.

This example illustrates with particular clarity a feature also exhibited by our other examples of subdoxastic states. Part of the reason we are inclined to insist that a subdoxastic state is not a belief is that if it were, we would be unable to say *what* belief it was. That is, if a subdoxastic state were taken to be a belief, there would be no sentence *p* such that inserting *p* in "*S* believes that _____" would express the fact that the subject was in the state. The problem is most noticeable when, as in our current examples, we contrast the belief that *p* with the subdoxastic state which stores or represents the information that *p*. Since the subdoxastic state differs so markedly from the belief both in its potential inferential connections and in the subject's potential access to it, we are disinclined to identify the subdoxastic state with the belief. But surely if a state representing the information that *p* is a belief at all, then it is the belief that *p*.

These reflections bring us perilously close to a cluster of questions I have been doing my utmost to sidestep. All this glib talk about a state representing information may get one wondering just what it might *mean* to say that a state "represents (the information (or fact) that) *p*." I do not think it would be unreasonable to duck the issue entirely in this paper. For though this talk of states representing facts is difficult to explicate in a philosophically tolerable way, it is surprisingly easy to master intuitively. Even the barest introduction to work in artificial intelligence and cognitive simulation quickly leaves one comfortable with attributions of content or representational status to the states of an information processing theory. And nothing I want to say here presupposes anything more than the ability to use these locutions as they are customarily used. Still, a hint of how I think such talk is to be analyzed may be welcome. On my view, saying that a state in an information processing system represents (the information (or fact) that) *p* is to say that the state bears some interesting resemblances to the belief that *p*. Generally the resemblances are with respect to the ordinary causes of the state, or some part of the inferential pattern of the state, or some other (non-inferential) effect the state may have (or some combination of the three). On this account, of course, the belief that *p* represents (the information (or fact) that) *p*; indeed, it is the prototype of such representation. Also, on my account, many different states can represent *p*, and can do so to varying degrees. There is no minimum degree of resemblance required for us to say that a state represents *p*. Rather,

the appropriateness of the content attribution depends on the particular needs and interests of the project at hand. For more on all this, along with a theory about what we are saying when we say a belief is the belief that *p,* see my *The Case Against Belief* (in preparation).

Our third example of a subdoxastic state, the Hess example, is quite parallel to the second. The subdoxastic state which serves to represent the information that the pupils in one photo are larger than those in the other is significantly less inferentially integrated with the body of a subject's beliefs than is the *belief* that the pupils in one photo are larger. As in the previous case, the paths leading to formation of the subdoxastic state are relatively few and involve only a restricted range of beliefs (if any), while almost any belief can play a direct role in an inference to the *belief* that the pupils in one photo are enlarged. Also, the subdoxastic state has a significantly more restricted set of possible inferences in which it may play a role. The belief that the pupils in the retouched photo are larger, along with the belief that, say, the earlobes in that photo are much enlarged, may lead to the belief that several facial features are enlarged. But one of Hess' subjects who has the *second* belief is unlikely to infer the belief that several facial features are enlarged (unless, of course, the subdoxastic state gives rise to the conscious *belief* that the pupils are enlarged). I think it is the inferential isolation of the state Hess' subjects are in that incline us to insist these are not to be counted as beliefs.

It is time to take stock of our discussion so far. I began by noting the distinction between beliefs and subdoxastic states, which is quite fundamental to our ordinary, pre-theoretic thinking about beliefs. With some examples of subdoxastic states on hand, we began to look for the basis of our intuitive distinction. My thesis has been that the distinction is drawn on the basis of two characteristics which beliefs exhibit and subdoxastic states do not: access to consciousness and inferential integration. In the course of our discussion we also noted, if only to then ignore, a further principle embedded in our pre-theoretic notion of belief, which ties together the concepts of belief and inference. The principle holds that inference (at least when viewed as a relation among psychological states) is a relation exclusively among beliefs. This principle, along with the features which divide beliefs from subdoxastic states, accounts for the location in traditional epistemology of non-inferential beliefs. For if the psychological states causally underlying a belief are not themselves beliefs, then, according to the principle, the process leading to the belief cannot be inference.

7. This principle that inference is a relation among beliefs deserves

a less cavalier look than has been accorded it so far. For the principle plays a central role in an argument advanced by Gilbert Harman aimed at undermining the belief-subdoxastic distinction. Harman notes that "a person determines how far away a perceived object is by means of cues involving overlapping of surfaces and texture gradients. It is natural, he continues, "to describe this as a matter of inference: given these cues the perceiver infers that objects are in those places. The relevant cues make it reasonable to suppose the objects are where the perceiver infers they are" ([7], p. 175). Harman anticipates the objection that talk of perceptual inference uses the word "inference" in a peculiar or special sense. He replies as follows:

> Contemporary psychologists tend to view a perceiver as an information processing mechanism, a kind of analogue computer. . . . Now it is natural to describe mechanical information-processing machines—like computers—as if they could calculate, figure something out, and infer conclusions. When the perceiver is conceived as an information processing device, it becomes natural to describe him in the same way. Having extended the application of 'inference' so that computers can be said to infer, it is natural to extend it so that perceivers are also said to infer. ([7], pp. 176–177)

The view that inference is involved in the formation of beliefs about our perceived environment is hardly novel. Descartes would have had no objection. But Descartes would have balked at Harman's view of the premises for such inferences. In the Cartesian tradition, the premises are beliefs about experience, about how things look to the subject. Harman rejects this position. "One problem is that some aspects of the way things appear is determined by inference. For example, reasoning involving overlap, texture gradients, and perspective figures in the apparent location of objects. Inference gets into the story before it is determined how things look" ([7], p. 180). Nor, on Harman's view, is there any more basic level of visual experience to serve as data. ". . . there does not seem to be any more basic level of visual experience not itself the product of inference, and used itself as data for inference to how things look" ([7], p. 181). Instead, Harman traces the path of inference all the way back to retinal stimulations, which, he holds, are the basic data for perceptual knowledge.

> I suggested that the data are the sensory stimulations that provide the input to the complex information processing system composed of brain and nervous system. Perceptual knowledge is based on inference from sensory stimulations. ([7], p. 185)

Finally, Harman claims that retinal stimulations and other sensory stimulations which serve as the basic data for perceptual knowledge are themselves *beliefs*.

> Does he [the perceiver] believe the data? He uses them in the way one uses beliefs in inference. So there is some reason to say that he believes them. ([7], p. 186)

Now note that if we go along with Harman, then the distinction between beliefs and subdoxastic states seems to vanish. For on Harman's view, both retinal stimulations and all of the various psychological states that may be inferred from them count as *beliefs*. If Harman is right, then our intuitive psychological theory has led us seriously astray; many states which intuition insists are not beliefs turn out to be beliefs nonetheless. Let us take a more careful look at just how Harman reaches this surprising conclusion.

I think the argument is best viewed as proceeding in two stages, the first defending the view that retinal stimulations and various other intuitively subdoxastic states are inferentially related to beliefs and to each other, and the second marshalling the principle that inference is a relation among beliefs. The claim of the first stage, in turn, follows from a pair of contentions: first, that it is appropriate to use the word "inference" in describing the processes leading to beliefs about our perceived surroundings; and second, that it is suitable to describe these processes as inferential "all the way down," to the level of retinal stimulations.

I am largely in agreement with the claim that there are inferential relations among (intuitively) subdoxastic states, and among subdoxastic states and beliefs. However, I think that Harman's defense of this view deflects attention from the important issue involved. In arguing that inference is involved in the process leading to the formation of beliefs about our perceptual environment, Harman stresses the "naturalness" of this use of the word "inference," a naturalness which is allegedly enhanced by the analogous use in talk of computers. It seems to me, however, that the intuitive naturalness of this use of "inference" is largely beside the point. What is important is not whether perceptual psychologists are using "inference" in some natural extension of its ordinary sense, but rather whether the phenomena they are describing are in important ways similar to more standard cases of inference. Similarly, the conjecture that inferential processes extend all the way down the causal chain to the level of retinal stimulations does not turn on the intuitive naturalness of calling these processes "inferential." What is at stake is whether these processes are in fact in fundamental ways similar to the process of inference

in standard cases. These are empirical questions, not to be settled by appeal to ordinary usage. What we need to know are the similarities and differences between the process leading from sensory stimulation to belief, and the process (ordinarily called "inference") leading from pre-existing beliefs to the formation of new beliefs. If the processes are sufficiently similar, the proposed extension of "inference" to cover the former is a reasonable one. At present neither perception nor more standard cases of inference are sufficiently well understood in detail to settle the question definitively. But recent efforts at modeling perception and at modeling inference and problem solving make it plausible that the two processes may be essentially congruent.[3] If this speculation proves true, it will be an empirical result of enormous importance.

The discovery that inference is involved in perception down to the level of sensory stimulation is not, of course, enough to show that intuitively subdoxastic states are beliefs. To get this conclusion we need the additional premise that inference is a relation among beliefs. Harman, as we saw in the last displayed quote above, endorses the premise and draws the conclusion. It is just here that my dissent comes in. Harman and I are agreed that inference likely relates states far removed from anything we would pre-theoretically call a belief. This fact, if it is a fact, puts considerable strain on our pre-theoretic views about beliefs. Harman would accommodate the strain by rejecting our intuitions on what sorts of states to count as beliefs. By contrast, I would accommodate the strain by simply rejecting the principle that inference is a relation exclusively among beliefs. By my lights, what the supposed discovery of inference in deep perceptual processes shows is not that intuitively subdoxastic states are actually beliefs, but rather that the domain of the inference relation includes other states in addition to beliefs.

There is something misleading in the account I have just given of the difference between Harman and myself. For it sounds as though the dispute is largely terminological. He wants to call subdoxastic states "beliefs" and preserve the principle that inference relates only beliefs; I want to renounce the principle and abjure the expanded extension of "belief." The issue, however, is not simply a terminological one. What is at stake is not whether we bloat the extension of "belief" to preserve a principle, but rather whether the old intuitive

[3]For problem solving see [12]; for some interesting work in inference, c.f. [1], ch.13; for an account of perception which stresses the parallels between perceptual representation and inference on the one hand, and the representation an inference among beliefs on the other hand, see [1], ch.8; also [11] and [13].

boundary between beliefs and subdoxastic states divides states which are in fact psychologically different. If it does, then following Harman's terminological proposal would simply make it necessary to coin some new term to denote the class of states we are now intuitively inclined to call "beliefs."

Harman nowhere addresses himself to the issue of the psychological importance of the belief-subdoxastic state distinction. However, in light of his proposed extension of "belief" to cover subdoxastic states and his failure to propose some other way of marking the boundary, I think we might plausibly attribute to him the view that the distinction is of small importance. Yet, as some of Harman's further remarks illustrate, this is not an easy position to espouse consistently. Thus consider the following dilemma. After defending the view that various (intuitively subdoxastic) states, including retinal stimulations, are beliefs because they are used "the way one uses beliefs in inference," Harman is in something of a quandary about just *what* beliefs they are. Just what belief are we to attribute to a perceiver, for example, when his retinal nerve is stimulated in a certain way? An obvious answer: the perceiver believes that his retinal nerve is being stimulated in that way. But, sensibly enough, Harman is unwilling to buy that answer. For he conjures the following objection:

> The typical perceiver knows little psychology and nothing about stimulations of his retinal nerve. So if the data needed for his inference includes claims about his sensory stimulation, he does not believe the data. ([7], p. 185)

Harman is plainly troubled by the objection, and his efforts to parry are a bit bizarre. The perceiver, he concedes, does not have any beliefs *about* sensory stimulations. Rather, the sensory stimulations themselves are the beliefs; it is they which serve as data for perceptual inference

> The data are not about sensory stimulations, they *are* sensory stimulations. . . .
>
> Does [the perceiver] believe the data? He uses them in the way one uses beliefs in inference. So there is some reason to say that he believes them. ([7], p. 186)

Well and good. The stimulations are beliefs. But beliefs that *what*?

> Sensory stimulations serve the perceiver as non-linguistic representations which cannot be easily put into words. . . .
>
> In order to be able to express in words the input data, we would have to know much more about the system of representation

> and its functioning than we now know. Even then it might be impossible to find linguistic representations equivalent to the non-linguistic representations constituted by sensory stimulation. ([7], p. 186)

Now it is pretty clear that part of this reply is simply a bluff. Harman suggests that we might be able to express the content of sensory stimulations if only we knew more about the workings of the system. But it seems pretty plain that there is no sentence which will serve to express the content of retinal stimulations *any better* than a sentence asserting that the retina is being stimulated in a certain way. Indeed, any sentence S which might be proposed as a candidate for expressing the content of retinal stimulations would face just the same difficulties as a sentence specifying how the retina is being stimulated. To see this, consider for a moment why we are unwilling to say of a perceiver whose retina is being stimulated in a certain way that he *believes* it is being stimulated in that way. Harman's reply is that the typical perceiver knows little psychology and nothing about the stimulation of his retinal nerve. But why are we so sure of this? The answer, I think, is that if we were to attribute beliefs about retinal stimulations to a naive perceiver, they would be a most peculiar species of belief, for they would be beliefs to which the perceiver had no access and which were largely inferentially isolated from the remainder of his beliefs. The same would be true if we instead attributed to the perceiver the belief that p, where p is an arbitrary candidate for expressing the content of a retinal stimulation.

Harman's other strategy for parrying the problem that we cannot find suitable content sentences for retinal stimulations is to retreat into mysticism. Even after we have a decent understanding of the functioning of the perceptual system, "it might be impossible to find linguistic representations equivalent to the non-linguistic representations constituted by sensory stimulations" ([7], p. 186). Their content is so mysterious that it is not expressible by words. It is tempting here to ponder whether it could make any sense to say that a state represents something, but there is no saying what it represents. I shall, however, resist the temptation. For it seems to me that by retreating to this doctrine of the unspeakable, Harman has conceded the principal point that I have been arguing. The states for whose contents it is impossible "to find linguistic representation" are just those that we have been calling subdoxastic. So Harman has conceded, however obliquely, that there is an important difference between beliefs and subdoxastic states. But if Harman and I agree on this central

point, serious differences remain in how to characterize the difference. For Harman, the difference turns on the linguistic expressibility of the contents of subdoxastic states. On my view, however, there is no particular problem in expressing the contents of subdoxastic states. Harman sees a problem because he insists that subdoxastic states are beliefs. But it is more than his intuition will tolerate to attribute to a subject the belief that *p*, where *p* expresses the content of the state. Thus he must postulate some mysterious unspeakable content. As I see it, the problem is not in finding a content sentence, but in failing to distinguish subdoxastic states from beliefs. Once the distinction has been recognized, we can assign quite ordinary contents to subdoxastic states without worrying that the subject fails to believe the content sentence.

We began this section with the concern that Harman's arguments appeared to show the distinction between beliefs and subdoxastic states is untenable. I think it fair to conclude that we have found the fault to lie not with the distinction but with the argument trying to undermine it.

8. Where does all this leave us? I believe we have a plausible case for the claim that the intuitive distinction between beliefs and subdoxastic states marks a real and psychologically interesting boundary. Moreover, it is a boundary that has been largely overlooked by contemporary work in cognitive simulation. The reasons for this neglect are not obvious. It would be my guess that, like Harman, many of those concerned with cognitive simulation have been so captivated with the promise of inferential accounts of the mechanisms underlying perception and thought that they have failed to note the rather special and largely isolated nature of the inferential processes between beliefs and subdoxastic states. Failure to take seriously the matter of access to consciousness likely has a less creditable explanation. Since the heyday of behaviorism, conscious awareness has had a bad name among many psychologists. And the attitude seems to persist even among those who have come to see behaviorism as a dead end.

If it is granted that the belief-subdoxastic state distinction is (at least *prima facie*) an important one, then we are left with an intriguing question. Why do inferential integration and access to consciousness run in tandem? Is this just an accident, or is there some underlying mechanism that accounts for both phenomena? My hunch is that it is more than an accident. But if I have inferred that view, the premises I have used are well hidden from conscious access.

REFERENCES

[1] Anderson, J. and Bower, G. *Human Associative Memory.* Washington: John Wiley and Sons, 1974.
[2] Austin, J. L. *Sense and Sensibilia.* New York: Oxford University Press, 1964.
[3] Chomsky, N. and Katz, J. "What the Linguist Is Talking About." *Journal* of Philosophy 71 (1974): 347–367.
[4] Dretske, F. *Seeing and Knowing.* Chicago: University of Chicago Press, 1969.
[5] Graves, C., Katz, J., et.al. "Tacit Knowledge." *Journal of Philosophy* 70 (1973): 318–330.
[6] Haber, R. N. and Hershenson, M. *The Psychology of Visual Perception.* New York: Holt, Rinehart, and Winston, 1973.
[7] Harman, G. *Thought.* Princeton: Princeton University Press, 1973.
[8] Hess, E. H. "Attitude and Pupil Size." *Scientific American* 212 (1965):
[9] Hess, E. H. "The Role of Pupil Size in Communication." *Scientific American* 233 (1975):
[10] Julesz, J. *The Cyclopean Eye.* New York: Academic Press, 1970.
[11] Moran, T. P. *The Symbolic Imagery Hypothesis: A Production System Model.* Ph.D. Dissertation, Carnegie-Mellon University, 1973.
[12] Newell, A. and Simon, H. *Human Problem Solving.* Englewood Cliffs: Prentice Hall, 1973.
[13] Pylyshyn, Z. "What The Mind's Eye Tells The Mind's Brain: A Critique on Mental Imagery." *Psychological Bulletin* 80 (1973):
[14] Stich, S. P. "What Every Speaker Knows." Philosophical Review 80 (1971): 476–496.
[15] Stich, S. P. "Competence and Indeterminacy." In *Testing Linguistic Hypotheses,* Cohen, D. and Wirth, J. R., eds. Washington: John Wiley, 1975.
[16] Winograd, T. "Understanding Natural Language." *Cognitive Psychology* 3 (1972): 1–191.