FROM EXPERIENCE TO BELIEF: MODES OF MENTAL REPRESENTATION

The mind is an intentional organ, and intentionality is, at bottom, a representational affair. This is true of sensory as well as cognitive phenomena, true of our experience of, as well as our judgments about, the world we live in. There are, nonetheless, important differences between perceptual and conceptual modes of representation. Perceptual beliefs, unlike perceptual experiences, are representations whose constituent elements--concepts--have their origin in learning. As a result, the beliefs which these concepts make possible are, unlike the experiences that give rise to them, capable of explaining the behavior of the individual in which they occur.

What follows is an elaboration of this theme. I begin by remarking on a fact about beliefs--who has them and when--that I find suggestive.

1 Learning.

The propositional attitudes, those inhabitants of the mind that involve a deployment of concepts, only seem to figure in the explanation of behaviors that are a product of learning. We do not invoke beliefs and desires, intentions and purposes, to explain why a person snores and sneezes, blushes and breathes, hiccoughs and belches, shivers in the cold and perspires in the heat. As we all know, these are things people do. They are not, like growing old or getting pregnant, things that happen to them. Nonetheless,

though they qualify, in this sense, as behavior, there are no explanations for these behaviors in terms of what people believe and want, intend and expect. Jane jerks her hand away from the hot burner, yes, but not because she thinks that by doing so she will avoid a nasty burn. She may, in fact, think this, but thinking this is not why she pulls her hand away. This is a reflex having nothing to do with her beliefs and desires. Beliefs, desires, intentions and purposes, are reserved for others kinds of behavior; they explain intentional or voluntary behavior, what are called actions, and such behavior, explicable in terms of reasons, is precisely the sort of thing we have learned to do or can, at least, learn not to do. Beliefs and desires explain behavior whose occurrence is sensitive to consequences.

Why is this so? Why are beliefs and desires--and, it seems, all the other content bearing internal states that function as reasons for doing something--located in (and, it seems, only in) animals capable of modifying their behavior in the light of what happens to them? Why, furthermore, do these reasons only appear, why do animals only start having beliefs, internal states with propositional content, when they start exhibiting behavior that is in some way the product of learning? Very young infants, even when they get what they want (food or attention) by sucking and crying, do not suck and cry because they think it will get them what they want. There are, often enough, reasons why an infant behaves in this way: it is thirsty, hungry, or in pain. But these are causes, not reasons. The reasons for behaving in this way (to get food or

comfort) are not--not yet anyway--the reasons why the baby behaves this way.[1] That comes later. Then, as we all know, children start yelling because they think it will get them the attention they want. Beliefs and desires only put in an appearance when there is something--an action, something modifiable by learning--for them to explain.

Since the propositional attitudes appear on the evolutionary and developmental scene when, and only when, learning occurs, when there appears the kind of behavior--voluntary or purposive--that propositional attitudes are invoked to explain, the suspicion is irresistible that the elements of these explanations--the beliefs and desires we invoke to explain behavior--have their origin in precisely those transactions, the learning processes, that gives rise to the behaviors needing explanation. Beliefs and desires, internal states with content, emerge as internal states with content, as (therefore) mental states, in the learning episodes that make possible the kind of behavior that such content is used to explain.

Such, at least, is how I read the clues. That is why I think we need a theory of content that features the learning process. We cannot look to biology for content, at least not the kind that will help you understand why I came to Houston this weekend. Obviously we do not inherit our beliefs. Nobody, I assume, will argue about that. But I think it equally clear that neither do we inherit the concepts needed to have the beliefs that explain what we do. They, too, are a product of individual learning. What we inherit in the

way of representational capacity is the capacity to experience the world, the capacity to perceive the objects and events about which we, through a process of learning, come to have beliefs.[2] But this is quite a different representational phenomenon.

In order to describe the representational differences between experience and belief and to simultaneously illustrate the role of learning in conceptual representation, I will describe, in greatly oversimplified terms, an idealized development, a process in which an animal acquires the capacity to hold beliefs about objects it already perceives. What I hope to illustrate with this story is the difference between modes of mental representation--a sensory mode that is innate and a conceptual mode that is acquired. To confuse these modes of representation, as theories of content tend to do, is to confuse what it takes to see an X with what it takes to recognize an X, to see (hence, believe) that it is an X.

2. Experience: Sensory Representations.

The hero of my story is Buster, a generic animal.

Buster lives in a place where there are furry, worm-like creatures called (naturally enough) furms. Furms come in all sizes and colors, but they all are furry (F) and they are all shaped like worms (W). Since they are the only creatures in Buster's habitat that are both furry and worm-shaped, the properties F and W suffice to identify an object as a furm in this environment.

Furms and their salient features are observable by Buster. By this I mean that Buster comes into the world with eyes and ears, a

sensitive nose and a discriminating palate. After a short period of maturation--a few days for animals like Buster--he is able to see the shape and color, the texture and size, the movement and orientation, of middle-sized objects like furms. Buster doesn't have to learn to see furms. Since they are all around him, all he has to do to see them is open his eyes.[3]

Though Buster is, from birth, equipped to see furms, he has no instinctive, no genetically programmed, reaction to them. They are not, in this respect, like bright lights, hot surfaces, and looming shadows, the sorts of objects to which Buster has an instinctive response. Buster doesn't withdraw or hide from, snarl or stare at, attack or bite, furms. Or, if he does, this is not to a furm as a furm, but as (say) an obstacle in his path, something casting a large shadow, or as something moving in his peripheral vision. Buster might, out of curiosity, sniff at a furm, poke at it, step on it, or watch it if it does something unusual, but a furm for Buster, initially at least, is much what eucalyptus trees, oranges, or Toyotas are to someone living in California: just one of those uncategorized things in the environment that one cannot help see when one's eyes are open but to which, prior to learning, one pays no particular attention.

What I am asking you to imagine, of course, is an animal who can and does experience objects of a particular kind (furms) but an animal who as yet has developed no capacity to respond to them as objects of a particular type. Buster does not yet have the concept FURM. Though he sees them, though he has experiences of them, and,

therefore, in this sense, internally represents furms (more of this in a moment), Buster does not yet have beliefs of the sort: that (what he sees) is a furm.[4] He represents them, but not yet in a way characteristic of belief, judgment, and knowledge. There is a gap between the sort of internal representation required to see furms and the kind of representation required to believe that they are furms, and Buster has not yet bridged that gap. If Buster is to believe of the furms he sees that they are furms, if he is to see them as furms, he has to develop a new way of representing them. Unlike the mechanisms underlying his perception of them, nature did not equip Buster to recognize furms--at least not as furms.

It is, perhaps, worth pointing out that though Buster is conscious, perceptually conscious, of furms (he sees them), there is another sense (easily confused with this sense) in which he is not conscious of them. He is not yet conscious of them as furms, as members of a single class. Indeed, he does not yet see furms as anything except, perhaps, as objects of some sort. He is conscious of furms, but not conscious that they are furms.

2. Belief: The Development of Cognitive Representations.

Let us suppose, therefore, that, out of idle curiosity, Buster sniffs at a large red furm and experiences a painful stinging sensation in his nose. Buster thereafter avoids large red furms. But the same thing happens again with a small green furm. In a short time we find Buster avoiding all furms. He no longer sniffs

at them. When they approach, he quickly retreats. Why?

Let me put this explanatory question in a concrete context so as to clarify just what we are seeking when we ask to have Buster's behavior explained. Imagine that we put Buster in an artificial situation in which there are fake furms (caterpillars or, if you like, mechanical furms), objects that are both furry and worm-like but not furms. Buster spots a fake furm approaching, and quickly withdraws. Why did Buster retreat? What is the explanation of Buster's behavior?

First Try: Buster withdrew because he saw a furm coming toward him.

This cannot be the right answer since Buster did not see a furm coming toward him. It wasn't a furm.

Second Try: Buster withdrew because he saw what looked like a furm (something that was both F and W) coming toward him.

This, too, can't be right. At least it cannot be the full explanation because prior to his painful encounters with furms, Buster saw what looked like furms coming toward him (in this case they actually were furms), objects that were both F and W, and Buster did not withdraw. He had exactly the same experience then, before learning, that he is having now, after learning, the sort of experience that constitutes his seeing an approaching furm. Now he withdraws; then he didn't. Now the experience triggers withdrawal movements; then it didn't. Why? What is the difference? That, surely, is the question we need answered in order to understand why Buster is withdrawing.

triggering cause of explanation

In attempting to explain Buster's behavior by describing what object he sees we are giving what I elsewhere (Dretske 1988) called a triggering cause of behavior. Normally, a triggering cause of behavior is some environmental object or event that is perceived (or the perceiving of it). An antelope, for instance, sees a hungry lion approaching. The antelope runs. The approach of the hungry lion (or the antelope's seeing the lion approaching) is a triggering cause of the antelope's flight. But neither event, neither the approach of the lion nor the antelope's perception of it, explains why the antelope runs. If you want to know why the antelope runs you have to discover what I called the structuring cause of this behavior. You have to find out, not just what the antelope sees, but how the antelope sees it, what it sees the lion as. If you assume that the antelope sees the approaching lion as a lion, or at least as danger of some sort, then, assuming normal motivations, you know why the antelope ran. Otherwise not. The same is true of Buster. What we seek when we seek an explanation of animal behavior, particularly when that behavior is deliberate or intentional, is not the cause of the animal's believing what he believes but (among other things) the belief it is caused to have. Saying that Buster sees what looks like a furm coming toward him tells one nothing about why Buster behaves the way he does. What we are seeking is the sort of explanation afforded by the belief that Buster is thereby caused to have. The approach of a furm now causes (triggering cause) Buster to withdraw (whereas formerly it did not) because it now produces, in Buster, a belief that it did

what are they seeing it as?

not earlier cause. The explanation we seek is an explanation in terms of this belief.

Third Try: Buster withdraws because he thinks it is a furm and he is afraid of furms. If, at this early stage of learning, this description sounds a bit too fancy for Buster, we can say that he withdraws because he thinks it, the approaching object, is one of those furry-worm-like (F-W) creatures that stings his nose (S) when it get too close. Call this, Buster's concept of a furm, the FWSer (furry-wormy-stinger) concept of furms. Buster runs because he thinks it is a FWSer and he doesn't want to get stung.

We have now entered new territory. We are now giving rational explanations of Buster's behavior, explanations in terms of what Buster thinks and what Buster wants. We may suppose that Buster wanted to avoid pain long before his painful encounters with the furms. This, we may suppose, is a long-standing, a more or less permanent, instinct or desire.[5] But we cannot suppose that his belief, the belief that completes this rational explanation, was available to Buster before his learning experience with the furms. Buster's capacity to have such a belief, his acquisition of the concept FURM (or FWSer), must have been acquired in his encounters with furms. All he had before his painful encounters with furms is a way of representing them in a sensory mode: he could see them. Now he has a way of representing them that explains why he reacts to them the way he does.

All this strikes me as fairly obvious. This is what concept formation is all about. A brief learning episode of this sort is

enough to credit an animal like Buster with a rudimentary concept, a concept essential for the holding of the sort of belief that will, when a relevant piece of behavior occurs, and in association with appropriate motivational states and other beliefs, explain the occurrence of certain aspects of subsequent behavior.

The details of this process, though interesting and important, are not relevant to the current project. They are the business of learning theory and developmental psychology. What I am concerned with is what such a process presupposes about the representational mechanisms mediating Buster's perception of, and his reaction to, furry worms. Clearly, if beliefs are to be understood as internal representations, then, as a result of such learning experiences, a different mode of representation must emerge, a way of representing objects that, unlike an experience of them, is available to provide rational explanations of behavior in terms of how these objects are being represented. What Buster has developed during his brief but painful encounter with furms is a way of representing them that goes beyond the sensory mechanisms he was given at birth, beyond what it takes to see furms. What he now has is a representation that, unlike his experience of furms, helps explain why he now avoids them. What Buster now has is furm-ish beliefs, not merely furm-ish experiences.

In order to better understand this new form of representation I think it is useful to say a little more about what I mean by a representation. I think of a representation in informational and functional terms: a representation (in the first instance[5]) is the

informational product of some device, organ or mechanism having an information providing function. A speedometer, for example, is an instrument that has the job, the function, of providing the driver with information about the speed of the car. When it is working the way it was designed to work, the way it is supposed to work, it does its job by producing various orientations of a movable pointer on a calibrated dial. These pointer positions carry (are supposed to carry) the information it is the job of this device to deliver. Since pointer positions carry, or are (given the function of the device) supposed to carry, information about car speed, they are representations of car speed.

We are all familiar with such conventional representations: maps, diagrams, instruments, gauges, pictures, and perhaps most importantly language. I call these conventional representations because the functions in question are dependant on the purposes and intentions of their designers and users. Change these intentions and purposes, or the way we use these artifacts, and the functions change. And as the functions change, the powers of representation change. All metals, including thumb tacks and paper clips, expand when heated, thereby carrying potentially useful information about temperature. Nonetheless, the only pieces of metal that represent the temperature are those (like the mercury in a thermometer) that service a device having the function of providing information about temperature. That is why thermometers can, but thumb tacks cannot, misrepresent room temperature.

If we assume, as I do, that there are, besides conventional

functions, natural functions, functions that in no way depend on our purposes and intentions, our beliefs and attitudes, the sorts of functions associated with bodily organs and mechanisms, then we have a class of natural representations. A natural representation is the information carrying product of any organ or mechanism whose information-providing function in no way depends on our intentions, purposes, or use of it. If, then, we think of the senses, as it seems natural to think of them, as organs having an information providing function, a function they derive from their evolutionary history, then the information-carrying product of these organs, the internal experiences they give rise to, are (just like a speedometer's pointer readings) representations, in this case natural (not conventional) representations, of the world they provide, or have the function of providing, information about.

Buster's experience of furms carries information about the texture, size, color, orientation, movement and shape of furms. If we assume that this experience is the product of a sensory system having the biological function of supplying such information, then Buster's experience of furms is a representation of furms. Since the functions underlying these representations are phylogenetic, functions the organs have in virtue of their membership in what Millikan calls a reproductive family having the right selectional history, we can, for convenience, call them P-representations.

Prior to the sort of learning described above, however, these P-representations had been put to no particular use. They did not do anything. When things are working right, these elements carry

Learning: recruitment of info for control

information about the size, shape, color, texture, movement and orientation of the large red furm on the right. They do the same for the small blue furm on the left. They do this for every furm Buster sees in normal conditions (perhaps with foveal vision), but the fact that these internal structures carry such information does not--not yet anyway--explain anything about Buster's activities. We have internal representations that have not yet been harnessed to do any useful work.

Learning of the sort described above changes this. It changes it by recruiting some of these information-carrying elements for control purposes. Shape and texture indicators take on a special significance. In particular, the F and W indicators are recruited for special work because of the vital information they carry. They are pressed into executive service because they (jointly) indicate the presence of furms, the sort of external object with which Buster's movements must be coordinated if he is to avoid being stung again. Buster's control circuits have been modified, additional centers of command installed, as a result of his experience with furms.

learning has created indicators for a new function

One way of thinking about this reorganization in Buster is to realize that learning has imposed an additional indicator function on elements that were, before learning, servicing a phylogenetic function. The F and W indicators have been given, as it were, a new job. Biologically speaking, the function of these elements is to provide information about the specific texture (furry) and shape (wormy) of objects. But now, since furry worms sting one's nose,

tokens of these two elements have jointly assumed the function of signalling the presence of a stinger. F+W has come to mean FWSer or FURM to Buster.

Let me illustrate this transition from an P-representation (an experience of furms) to this new mode of representation, a mode I will call O-representation (O for ontogenetic) with a familiar sort of artifact. A manufactured instrument, a pressure gauge, has the function of registering external air pressure. Information about air pressure is displayed on a dial, suitably calibrated in pounds-per-square-inch, across which a pointer moves. Positions of the pointer represent the air pressure because the device has the job of providing this information by means of the pointer positions. Pointer positions are therefore representations. We can think of the instrument as having the (conventional) phylogenetic function of providing information about air pressure since it is its being an instance of a certain manufactured type, created with certain intentions and purposes, from which individual instruments inherit their information-providing function. Even if this instrument is never used as a pressure gauge, that is what it is.

Suppose, now, understanding the way air pressure depends on altitude, and needing something to inform us of altitude, we take one of these gauges and use it as an altimeter. We introduce a few adjustments to compensate for divergences between air pressure and altitude, recalibrate the dial in feet-above-sea-level, and install the gauge in the cockpit of our aircraft. Voila! A device whose phylogenetic function is described by calling it a pressure gauge

becomes an altimeter. States that served a phylogenetic function by indicating pressure are recruited to do another job--to indicate altitude. The internal components whose (phylogenetic) function it is to register changes in pressure still do their job, of course. The behavior of these components is now, however, at the service of a system whose acquired function is to indicate altitude above sea level. We have, as it were, imposed an ontogenetic function, the job of indicating altitude, upon a system which has a different (though usable) phylogenetic function, the job of registering air pressure. I like to think of this device as sensing air pressure but having beliefs about altitude.

The analogy is just an analogy, of course, since, in the case of artifacts, the functions in question, both the phylogenetic and the ontogenetic functions, are creatures of our (designers and users) intentions and purposes. Take us away and the functions, the representations, disappear; one is left with a piece of metal being caused to expand and contract by changes in air pressure. It is, intrinsically, and on its own, neither a pressure gauge nor an altimeter. We make it so. But we do not make it so with living systems. The functions, both those (phylogenetic) constituting the representational powers of sense experience and those (ontogenetic) underlying conceptual phenomena, are natural functions. There is nothing conventional about Buster's experience of furms or, after learning, his beliefs about them.

Buster's experience with furms has imposed an additional function on the elements (the F and W indicators) that were, before

Before learning Buster could not misrepresent

What it does rep defined by Buster's history of how F & W was recruited.

learning, only assisting in the performance of a phylogenetic indicator function. Before learning, F+W did not--indeed, could not--misrepresent an object as a furm. They could not make this kind of mistake because it wasn't their job to carry this kind of information. Learning has changed that. Just as the imposition of an additional function on the pressure gauge has made the pointer "say" something it never said before--hence made it susceptible of saying something wrong, of misrepresenting altitude--learning has made the F+W elements "say" something they didn't say before. It has made them represent the furry worms as furms, as furry wormy stingers, and it has done so by giving these elements the function of registering the presence of such undesirable creatures.

When our pressure gauge left the factory, it did not have the job of registering altitude. It was not an altimeter. It did not, therefore, misrepresent altitude even when it malfunctioned. Its power to represent, and thereby misrepresent, altitude awaited its assignment of this information-carrying function. The same is true of Buster's perceptual and cognitive mechanisms. When Buster left the factory, nothing in him had the job of registering the presence of furms. The senses, we are assuming, had the job of registering the texture and shape of objects--including, of course, the texture and shape of furms. But nothing had the phylogenetic function of indicating that the furms were furms. This should be obvious from the way I described Buster's condition. If it isn't, we may easily suppose that furms have just invaded Buster's habitat for the first time. These are the first furms he has ever seen. Nor have any of

Buster's ancestors ever seen a furm. Buster encountering furms may be like me encountering golf balls for the first time. Unless I am mistaken about my ancestors, none of them ever saw a golf ball. So my capacity to internally represent golf balls as golf balls, my capacity to believe that something is a golf ball, could hardly be innate. It is not the sort of capacity that we can look to biology to understand. It is an acquired power of representation. And the same is true of Buster. His ability to represent furms as furms, his capacity to believe of the furms he sees that they are furms, is not something he could have inherited from his ancestors.

What I inherited from my ancestors is not my concept of a golf ball, not my capacity to represent objects as golf balls, but the resources of internally representing the color, shape, size and distinctive dimples of golf balls, the properties that enable me now, once I have the concept of a golf ball, to represent them as (to believe that they are) golf balls when I see them. The same is true of Buster. What he got from his ancestors is the power to represent the shape and texture of furms, the properties that, once he got stung by them, would enable him to recognize them the next time he saw them.

Buster's O-representation of furms, though it is activated by his perception of furms, is a mode of representation that, in one sense, goes beyond his sensory representation. Buster's experience of furms is a representation (among other things) of F and W, the texture and shape of the furms he sees. His visual experience of furms does not include a representation of S, their power to sting

his nose. His developing concept of a furm, however, includes not just the visually experienced qualities F and W, but, in addition, S. Buster's visual experience of furms, though sufficiently rich in information (in his natural habitat) for identifying furms, does not exhaust his concept of a furm. His experience represents the F-ness and W-ness of the furms he sees, but his belief goes beyond this. It represents the furms he sees, not merely as furry worms, but as furms--furry-wormy-stingers. Buster's visual experience provides the information needed to identify furms, but it does not include all that is included in the belief that, after learning, this experience give rise to.

3. Belief and Behavior.

When we move from a P-representation to an O-representation, from Buster's experience (perception) of a furm to his belief that it is a furm, we move to representations that explain why Buster behaves the way he does. What it is that gives O-representations this explanatory power is their participation in the process, the learning process, in which the control circuits responsible for behavior are reconfigured. O-representations are created in the very process in which is formed the underlying structure of future behavior. That is why such representations, unlike inherited P-representations, can, qua representations, explain behavior.

Let me describe, briefly, the dynamics of this process and what it is about both the representations, and the behavior they explain, that makes this possible.'

To understand the way beliefs, qua beliefs, explain behavior, a representationalist (such as myself) must understand the way a representation, qua representation, explains behavior. This, in turn, means that we must understand how the fact that some physical structure is a representation can explain why it does what it does. Every physical structure, including those that are representations, have a variety of properties, some intrinsic, some extrinsic. Its being a representation is, on almost everyone's account of these matters, including mine, an extrinsic fact about it. What makes the events occurring in Buster's head represent the furms in his environment is the fact that they have a certain information-carrying function, whether phylogenetic or ontogenetic, and having this kind of function is an extrinsic (relational) fact about the events occurring in Buster's brain. We can change these relations (as Putnam imagined in his Twin Earth example) without changing anything in Buster. This being so, in order to understand how representations, qua representations, explain Buster's behavior, we must understand how extrinsic facts about Buster's internal states, those extrinsic facts that make them representations, help explain his withdrawal from furms.

This is a task that dooms almost all representational accounts of the propositional attitudes. To put it in Colin McGinn's (1989) terms, causality is a local affair; whether (and, therefore, how) something is represented is not a local affair; therefore, how we represent things is causally irrelevant to what we do. Internal representations, being physical states, may (as Davidson argues),

cause movements of the limbs, but they function in this way, not qua representations, but qua physical-chemical states of the brain. The mental, qua mental, is irrelevant--epiphenomenal.

Only by understanding the functions underlying representation as ontogenetic functions is it possible, I believe, to circumvent this objection.[8] For ontogenetic representations are created in the developmental processes that reconfigure the control circuits of the system in which they occur. As a result, they are directly involved, qua information-carrying states, in the production of any future behavior that results from that reconfiguration. Why does Buster withdraw when he sees a furry worm approaching? Because, as we common folk like to say, he thinks it is a furm, something that will sting his nose if it gets too close, and he doesn't want that to happen. According to a representational account, this means that there is, inside Buster, a representation of the furm as a furm, as a beast that will sting one's nose if it gets too close. Furthermore, this fact is made causally relevant to Buster's hasty retreat because Buster's representing furms in this way is his occupation of an internal state that was pressed into service in the control of behavior, including withdrawal behavior, because it carried this vital piece of information. Talking about what Buster believes by way of explaining what Buster is doing is just a way of describing the information-carrying properties of the controlling internal states--what information they are carrying, or supposed to be carrying, that led to their current influence on output.

Fred Dretske
Stanford University

ENDNOTES

1. There are interesting and relevant connections between "reasons why" (an explanation) and "reasons for" (rationalizations), but I do not have the time to discuss them here. I hope my use of the contrast is clear enough.

2. We also inherit whatever biological resources are required to develop these conceptual capacities, but this goes without saying. If we think of the concept of X as the capacity to hold beliefs of the form: so-and-so is X, then we inherit the capacity to acquire such capacities. This, though, is true of any capacity we might acquire--e.g., playing the piano.

3. Opening his eyes, and allowing the receptors and associated mechanisms to be stimulated during the first few days may be part of total maturation process that enables Buster to see furms. I set aside these complications as irrelevant to present purposes. What is relevant to present purposes is that this maturational process is quite different than the sort of learning process required for Buster to recognize furms.

4. Though if we assume Buster already had a concept, he could have a belief <u>about</u> (i.e., <u>of</u>) a furm that it was (say) blue.

5. There is a big difference between an instinct (to avoid X) and a desire (to avoid X), and the difference is relevant to explaining behavior, but I will not worry about these details here.

6. I add this qualification merely to register a point about my theory of representation. Some representations (what I later call O-representations) are merely information-carrying types whose tokens, because of the information they carry, have been given (or acquired) a special causal role. Since this is a technical point, I won't return to it again in this paper.

7. This is a superficial summary of the account I give in <u>Explaining Behavior</u>.

8. It is for this reason that I think anyone (Millikan? Papineau?) looking to biology as a (if not <u>the</u>) source of intentional content, especially if this content is supposed to be relevant to explaining behavior, are barking up the wrong tree. Biology may supply the functions that constitute P-representations (i.e., experiences) but not the functions needed for the kind of representations (cognitive or conceptual) that explain behavior.