

Epistemological Holism and Semantic Holism

William Cornwell
University of Connecticut

ABSTRACT: This paper draws upon the works of Wilfred Sellars, Jerry Fodor, and Ruth Millikan to argue against epistemological holism and conceptual holism. In the first section, I contend that contrary to confirmation holism, there are individual beliefs (“basic beliefs”) that receive nondoxastic/noninferential warrant. In the earliest stages of cognitive development modular processes produce basic beliefs about how things are. The disadvantage of this type of basic belief is that the person may possess information that should have defeated it but that was not taken into account by the module. For this reason, at more advanced stages of cognitive development basic beliefs concern how things appear to be. These appearance beliefs are not formed holistically but should be checked against background beliefs before the person infers how things are. In the second section I argue against functional role semantics / inferential role semantics. Championing teleosemantics, I argue that many concepts’ meanings are not determined by the meanings of other concepts. Rather, many concepts are skills of knowing how to identify of what the concept is. These skills can be developed independently of other beliefs or skills and are in an important sense theory-neutral.

Epistemic Warrant Without Holism

I am invited into your home for the first time. As I enter the dining room, I see a fruit bowl with two peaches, which naturally leads me to think that there are two peaches near me. This belief is typical of the sorts of beliefs that many foundationalists think of as being at the foundations of knowledge, now that most foundationalists have given up on there being foundational beliefs about sense-data and so forth. It seems as if this belief is formed directly on the basis of experience and has no inferential warrant – in short, that this belief is foundational but corrigible.

Consider a second case just like the first except that the belief is false because rather than two peaches being in the bowl, there are two expertly crafted waxen peaches. In this case, it seems that the foundational belief that there are two peaches in the bowl is warranted – it was reasonable to believe that there were peaches in the bowl – but false.

Consider a third and final case, which is just like the second one except that I also know that you work for a company that makes waxen fruit. In this case, most people would agree that I would not be warranted in believing that those were peaches in the bowl, even though they looked like peaches. Here, most of us are inclined to say, we have a case where prior beliefs (about you working for a company that makes waxen fruit) should block... the *inference* that these are peaches? But if there is an inference in case three, then there is an inference in case one which I did not block, yet we were supposing in case one that the belief that there were peaches in the bowl was foundational, hence noninferential. If the belief in case one was inferred from other beliefs, what were these other beliefs? If it is foundational, then why isn't the belief in case three also a warranted foundational belief? This is the foundationalist conundrum. I will offer a solution that is in the foundationalist spirit of single beliefs being able to receive nondoxastic warrant (thus my theory is not confirmation holism), but it is unclear that my theory is foundationalist.

To use a problematic foundationalist metaphor, the solution is to move the level of "foundational" judgments for the mature thinker in some instances one level down to judgments about how things appear, which I will call "appearance beliefs."¹ Thus in the above case my "foundational" belief is not that there are two peaches in the bowl but that it appears that there are two peaches in the bowl. In cases one and two, I properly inferred that I saw a peach from this appearance belief and my other beliefs, but in case three I should have inferred that I saw a waxen peach. I am suggesting that in mature cognition, we begin with appearance beliefs about seemings and make inferences from them about what is. These "foundational" beliefs about how things appear to be are fallible, mandatory, and are the basis of further inferences. Absent reasons to the contrary, when I believe that I seem to see a peach I should conclude that I see a peach; absent reasons to the contrary, when I believe that I seem to have an aching joint, I should conclude that I have an aching joint (and not, for instance, a phantom pain), etc. In short, absent defeasibility conditions, people do and should assume that things are as they appear to be.

As I mentioned above, these nondoxastically warranted beliefs are corrigible, so they are not, according to traditional foundationalist theories, foundational beliefs. Yet they are similar to foundational beliefs in one important sense: each has nondoxastic warrant. Thus I propose to borrow a term from foundationalist theories without committing myself to foundationalism and to call these beliefs "basic."

Presumably the inferred belief about how things are usually wipes out the appearance belief without us noticing that this has happened. For instance, the belief that I seem to see a peach leads me to believe that I see a peach, and after the inference is made, the appearance belief

is forgotten or discarded. Thus at the conscious level, we generally think about how things are and not about how they appear to be. The erasure of appearance beliefs makes a certain amount of sense. If I were challenged about whether I really saw a peach, it would not be difficult to reconstruct that I had believed that I had seemed to see a peach, so not having stored in memory this belief about how things had seemed to be entails little epistemic risk. What is more, the phenomenon of belief perseverance, in which beliefs persist even after their inferential bases have been destroyed, provides independent evidence that we do not keep track of the bases of our inferred conclusions. Perhaps this is because the resources of human memory are limited, and it is an efficient strategy to maintain conclusions while erasing their inferential bases (Harman 1997, 334-5).

The inferential process from how things appear to be to how things are generally will involve less than global considerations. Here speed trumps thoroughness, but this is OK, because more thorough coherence checks can be done after the initial inference. If I am alone on a sailboat at sea and hear a roar, as if a lion were just behind me, my lion-identification mechanism will kick in and form the belief that I seem to hear a lion. Despite the implausibility of there being a lion nearby, I might initially and subconsciously infer that there is one (and be very scared), because at first pass there is no plausible alternative explanation for this roar. Yet, if I turn around, see that there is no lion in the skiff with me, and consider how there could not be a lion at sea, I will cast about for an alternative explanation of what caused the sound, perhaps inferring that I heard a sea creature that sounds like a lion. The first-pass inference from the basic belief and some background beliefs usually falls short of holistic standards but is good enough, and when it is not good enough (as in the lion example), ideally more global considerations kick in later so that incoherence is temporary.

One advantage of a theory such as this one is that it can explain our beliefs in cases such as the Müller-Lyer illusion, in which even after the person is apprised of the equality of the lines, the belief that the lines seem to be of unequal length persists. The explanation is that once the person understands the illusion, he or she blocks the inference from the lines seeming to be of unequal length to the lines being of unequal length. As Fodor has argued, that beliefs about how things seem to be persist even after we know how things are suggests that many beliefs about how things seem to be are mandatory and encapsulated.

If this theory is true, in many cases we first form beliefs about appearances and then decide based on these beliefs what conclusions about how things are would best cohere with our other beliefs, at least a regional level. This does not mean that starting at birth we build up our picture of the world from appearance beliefs; after all, appearance beliefs

come on the scene long after beliefs about how things are. We teach a baby that this is a rattle and that is Mommy; we do not teach the baby that this seems to be a rattle and that seems to be Mommy. Learning about how things seem to be is parasitic on prior understandings of how things are. In the earliest stages of life, the belief-formation mechanisms fix basic beliefs only about how things are with no inferential input.² At this point, the child's mental processing operates to produce basic beliefs about how things are. At first, children's mistaken judgments are corrected or are noted by the children as they compare their judgments and behaviors to those of more cognitively advanced people, and the existence of mistakes is accepted with no need for further explanation. With more cognitive maturity, the child wants to understand why she made a mistake, so we introduce to her talk about how things seemed to be in order to explain some perceptual errors, for it turns out that some false perceptual beliefs are the understandable result of one thing (or kind) having appeared to be something else.³ Eventually the child understands this type of explanation and tacitly grasps under what conditions, say, the peach-identification module can be triggered by stimuli that are not peaches but, as we say, resemble or appear to be peaches. This explanation is extroverted – it focuses upon the idea that some things *in the mind-independent world* resemble each other. At an even more advanced stage of cognitive maturity, the child learns that often when mistaking something that is not a peach for a peach she was in a state qualitatively identical to or similar to that of when she has correctly identified a peach, even though she unknowingly perceived something different. For instance, the state of seeing a real peach from a distance and the state of seeing a waxen peach from a distance can be qualitatively the same or similar. In this case the child understands to some extent that what makes two things resemble each other is that they provoke the same perceptual response. (When the child grows up to be a philosopher, she recognizes the Lockean point which he expressed in terms of secondary qualities, namely, that what will provoke sameness of cognitive response is relative to the constitution of the perceiver and is not solely an “objective” property.) This explanation is more introverted and sophisticated, involving as it does a theory of the mind. In both explanations it appears to her that she is seeing a peach. Understanding the simpler or the more advanced explanation of error opens the door for the child to begin to make basic judgments about how things appear to be and to infer from a basic belief and some background beliefs how things are, which involves understanding under what circumstances, including psychological circumstances, certain types of errors are likely to occur – that is, recognizing deviations from the cognitive design plan – hence in those circumstances blocking the inference from “It seems that X” to “X.” In short, children advance from making basic judgments about

how things are to using collateral information (again, often short of global standards) to sort out whether things are as the child believes that they appear to be.

There are two routes to the fixation of a belief about how things appear to be. On the one hand, such a belief can come about from modifying a prior module responsible for issuing beliefs about how things are. For instance, my module that had been responsible for identifying peaches could be altered to sometimes issue instead judgments that there appear to be peaches. This is not a radical change, but it has the advantage of, as it were, making the modular processing sometimes stop at an earlier stage so that more holistic considerations can influence our perceptual judgments about how things are. In such a case, the proper function of the module is to form the belief that, say, there appear to be peaches only if there are peaches. Hence, even if there are waxen peaches on the table that appear to be peaches, the module should not produce the belief that there appear to be peaches on the table. If it does produce such a belief, the holistic checker should defeat the inference to the belief that there are peaches on the table, which is the advantage of having an additional level of cognitive processing. I will call such beliefs about appearances “first-order appearance beliefs.”

The other route to fixation of beliefs about how things appear to be is, from a phenomenological standpoint, the more familiar route for forming appearance beliefs: introspection. Whereas the first method requires only looking squarely at the world, the second method requires a form of reflection. On this approach, I can form the true belief that it seems that there are peaches in front of me by noticing that I am in the same type of internal state that has been part of the process whose proper function is to produce the true belief that there are peaches in front of me. This has some affinity to a two-factor theory of mental content. On the one hand, to be thinking about peaches requires the right past or present external conditions – this is one moral of Twin Earth cases. On the other hand, there will be certain types of internal states that have been selected as vehicles for external content, and insofar as the proper functions of those internal states have been to enable the person to recognize X, being in those internal states is seeming to recognize X, regardless of whether the person is recognizing X. If I unknowingly am whisked away to Twin Earth and see XYZ, then I seem to see water, because I am having the same type of internal state that in the past had been part of a process whose proper function was to lead me to judge that there was water. It is consistent with this account that an internal state could be, from a subjective standpoint, ambiguous, because that state has been selected in some conditions to carry one type of content and in other conditions to carry a different type of content. For instance, in some circumstances, the internal state of a person when seeing a peach can be identical to the

internal state of the person when seeing a waxen peach, and if, say, the former state (with the proper external circumstances) has had as its proper function to enable the person to identify a peach, then both states constitute the condition of seeming to see a peach, and understanding this point is what makes certain types of mistakes intelligible to us. This is similar to how a token of “bank” can be ambiguous from an interpretive standpoint, even though its meaning may be fixed by its external relations, namely, by considerations of to which reproductively established family it belongs. Hence in reflective consciousness a person can use an “introspective” module that is responsible for observing his own internal states in order to determine how things appear to be. I will call such beliefs about appearances “second-order appearance beliefs.” It might be that there is a distinct introspective module for recognizing that it seems to be cold and a different introspective module for recognizing that it appears that the apple is green, or it might be that there is a single introspective module, a central processing unit of sorts, that inductively recognizes the linkage between certain types of states and certain types of true judgments. We can ignore these details.

To these two types of appearance beliefs we can contrast the state of it appearing that so-and-so. For it to appear that so-and-so to person P is for P to be in an internal state that is identical to or resembles an internal state whose proper function is to produce the true belief that so-and-so. This internal state is not necessarily doxastic; it can appear to P that so-and-so without P thinking that so-and-so or thinking that it appears that so-and-so – there is nothing “given” about being appeared to. If there are such *things* as appearances, as compared to states of being appeared to, we can think of them either as being identical to these internal states or as being elements of them, although, as stressed earlier, the inherent properties of the appearances do not determine of what the appearances are. Thus we can distinguish appearances (or being appeared to), first-order appearance beliefs, and second-order appearance beliefs.

The two takeaway points for epistemology are as follows. First, if one would be warranted in the basic belief that there are two peaches in the bowl, then one would be warranted in the basic belief that there appear to be two peaches in the bowl. Hence when the module is modified to produce sometimes the latter belief instead of the former, the latter belief is supposed to map the same state of the world, namely, the two peaches being in the bowl. Yet, in one crucial respect the first order appearance belief functions differently from the belief that there are two peaches in the bowl, in that no appearance belief directly guides behavior. Rather, an appearance belief is used in the coherence checking process to lead the cognitive system to produce a belief about how things are in the world (and, in unusual cases, to alter the person’s background beliefs). It is the beliefs about how things are that guide behavior. The

difference between the first-order appearance belief and the belief about how things are concerns how they function in the person's cognitive economy, not what they map. Put another way, in first-order appearance beliefs, "It appears that..." is a functional operator upon regular indicative sentences and the indicative sentences carry all of the content. Perhaps these "appearance beliefs" are misnamed. As long as their nature and function are understood properly, it does not matter what we call them.

Secondly, if on the basis of introspective evidence one forms the belief that there appear to be two peaches in the bowl, this belief is supposed to map directly a state that is identical to or resembles an internal state that has as a proper function to bear the content that there are two peaches in the bowl. These states constitute the belief's truth conditions. Such a belief indirectly is supposed to map that there are two peaches in the bowl, or, if not map this fact, provide defeasible evidence for it. Ignoring this secondary content, first-order and second-order appearance beliefs do not map the same type of events and do not have the same truth-conditions, but the appearance beliefs have similar roles in terms of a person's inferential processing. In both cases, the beliefs make a prima facie case for there being two peaches in the bowl but the beliefs are not to be acted upon – they are to be used as part of the evidential basis for further inference. This is why, for simplicity's sake, I have treated and will continue to treat them both as beliefs that two peaches appear to be in the bowl, despite the fact that they have different direct content. Perhaps beliefs with different contents should not receive the same appellation, but no immediate nominative replacement suggests itself to me.

If the preceding account is right, then reflectively coming to see how things appear to be is coming to understand that a certain type of internal state sometimes has had as a proper function leading one to recognize how things are. As we grow up, we come to realize that we sometimes are wrong and that this internal state is not a perfect indicator of external conditions. We can improve our odds of being right about how the world is by *first* determining how things seem to be and then considering whether how things appear to be is likely to be how things are, given other information that we possess. The idea is that there are various types of beliefs that link appearances and reality, such as "Waxen fruit appear to be fruit," and that these beliefs can come into play in appropriate circumstances (as in example #3). As for first-order appearance beliefs, here too the ability to introduce more global considerations before fixing the belief that, say, there are two peaches in the bowl is likely to increase the reliability of such beliefs, because they no longer are produced solely by a modular process that by its nature has no access to information that can defeat the module's conclusion.

In general the inference from believing that one seems to recognize X to believing that one recognizes X is fast and unimpeded, and, as I mentioned before, involves “erasing” the former belief, especially if the former belief is a first-order appearance belief. (Second-order appearance beliefs are more likely to be conscious, hence to be remembered.) The inference will occur unless there are fairly obvious reasons for it not to occur. After all, the fixation of perceptual beliefs should be quick, so this inference is treated as innocent until proven guilty. Yet, there are times when our beliefs about how things seem to be run headlong against the contextually sensitive background of beliefs (“the frame”) that structures our expectations. This can cause us literally to start, as when it appears that someone who is supposed to be on the other side of the world walks into the room. Here the inference the person should draw from the appearance belief depends upon the relative strength of warrant of other relevant beliefs, such as that the person is overseas, that nobody else so resembles this person, and so on. Generally speaking, we presume a high degree of nondoxastic warrant for basic beliefs such as that it appears at that time that Mr. X is walking in the room, so that belief usually is taken as nonnegotiable, although as the person approaches, he may no longer appear to be Mr. X, and the threatened crisis of epistemic incoherence may resolve itself easily. In other cases, we might await further information (e.g., hearing the person speak) before deciding what inference to draw about how things are.

I will make a few final points to head off potential misunderstandings. First, the claim of encapsulation of the modules responsible for first-order appearance beliefs needs qualification. Whether that module will produce the belief that it appears that... may depend not only upon the input to the module (some inputs should almost always activate the module, but other inputs are borderline cases) but also on the extent to which the module is activated or quieted by other beliefs or linguistic stimulation. For instance, suppose that I see someone indistinct far off in the distance, and I have no belief about who it is or may be. If someone says, “That’s Mary,” then suddenly I may believe that it appears to me to be Mary. The mere mention of Mary’s name might be enough to make the Mary-detecting module become more sensitive in issuing the belief that it appears to be Mary. Of course, this module is imperfect, and under these circumstances it may be overactive relative to modules for identifying other people, so that I become too susceptible to forming beliefs that I seem to see Mary. With more cognitive maturity, one learns to check these tendencies by, say, thinking to oneself, “It may appear to be Mary, but at this distance I can’t be sure.” In any case, the attenuation of activation strengths of various belief-formation modules, which presumably is also what is involved in focusing attention, can in certain ways mimic the effects of a coherence-check without being identical to it.

After all, the mere mention of Mary's name not even in the context of a statement about Mary (as in "Mary!" or "Is that Mary?") might increase the activation strength of this module without any belief about Mary having done so. A similar phenomenon can occur with concepts associated as co-occurring.⁴ For instance, we often associate a person with his or her spouse, so if I see two people together at a distance and identify one who has a spouse, then this can make the module for identifying the spouse more sensitive, making me more likely to believe that I seem to see the spouse. This ability to attenuate the sensitivity of a concept in light of the activity of other concepts – to link concepts – strikes me as psychologically plausible and on balance epistemically desirable.

Secondly, I want to emphasize that the operations of at least some modular systems are supposed to be synchronously fixed (except in the way discussed above), not diachronically fixed. Indeed, it is essential to the theory I am propounding that a person (a) can *develop* (or be trained to acquire) a fallible mechanism for identifying how things are – these mechanisms are not innate – and (b) then can *transform* the mechanism so that some of the time its output is first-order appearance beliefs that no longer guide behavior. Perhaps some modular belief-fixation systems are hard-wired, but many are not. After all, we do *learn*. As I mentioned earlier, whether the systems responsible for producing sensations are hard-wired and exactly what role sensations play in the fixation of perceptual beliefs are questions that we can ignore.

The inference from a belief about how things seem to be to a belief about how things are is parasitic upon and is a refinement of the earlier ability to form true judgments about how things are, and there is no guarantee that every ability of the latter sort will be replaced by an ability of the former sort. In all likelihood some, maybe even most, of a mature adult's basic beliefs are of how things are rather than how things seem to be, and the former beliefs also will be the result of modular processes that can be affected by other beliefs in the way discussed earlier, a way that falls short of inference. My point is that insofar as we think that a person is culpable for errors in perceptual beliefs about how things are that result from not taking into account collateral information, we are claiming that the belief was not basic or at least should not have been basic. There is a trade-off between speed and accuracy, so that faster but less accurate are basic beliefs about how things are than inferences from how things appear to be to how things are. Yet the latter inferences are fast when there are no obvious beliefs to undermine the inference from the belief "It appears that X" to the belief "X."

Concepts Without Holism

Inferential role semantics (or “functional role semantics”) threatens my theory of basic beliefs. Inferential role semanticists think that a concept is to a meaningful intensional state as a meaningful sub-sentential term is to a meaningful sentence. Intensional states include beliefs (which are, for our epistemological purposes, paramount) but also desires, musings, wonderings, etc. There are many varieties of inferential role semantics, but a rudimentary formulation of it that should be enough for our purposes is that the meaning of concept X depends upon the meanings of lots of other concepts that participate in some of the same judgments as X, or, to put it another way, that the meaning of X depends upon its inferential connections to lots of other concepts.⁵ Whether inferential role semantics should be developed on molecular or holistic lines will not affect my argument.

Inferential role semantics leads to something like confirmation holism. Confirmation holism can be given a fairly trivial reading, in which one’s entire empirical “theory” (in the broadest sense of being all of one’s beliefs) confronts all of one’s experience. It is hard to see why any empiricist would want to deny such a point, at least as an epistemological ideal. The more contentious claim of confirmation holism is that there is no way to confront a particular belief with a particular experience, because a precondition of the truth of any particular belief is the truth of a broader theory of which it is a part, so no part of that broader theory could confront reality or be tested against reality independently of that theory. (The broader theory could be the total of one’s beliefs or could be something more regional, such as physics.) The confirmation holist admits that there can be atomic facts but thinks that no meaningful belief (hence no belief) about an atomic fact can exist unless that belief belongs to a broader theory which bestows upon the belief its inferential role, so there can be no true belief without the other beliefs that constitute the theory also being true. This means that the theory stands or falls as a whole, and there is no part of that theory can be tested independently against experience. If this is so, not only can there be no foundational beliefs but there also can be no basic beliefs. Sometimes this objection is expressed in terms of every judgment being “theory-laden,” as compared to there being “theory-neutral” basic or foundational beliefs.

I will sketch an alternative, teleosemantic theory that does allow many concepts to be tested independently against experience, although there are some concepts that stand or fall together. I also want to argue that once “theory-laden” and “theory-neutral” are disambiguated, we will see that my theory acknowledges that in a certain sense every belief is theory-laden but that in another sense some beliefs are theory-neutral.

Following Ruth Garrett Millikan, I think of concepts as abilities to track things or kinds (both of which she calls “substances”) through space

and time, gathering information about them as we go. We can individuate a concept in terms of what is tracked, so that two people have concepts of the same type if and only if the concepts track the same thing. Of course, there can be differences in how things are tracked by different people. Using Millikan's terminology, people in such a case have conceptions of different types but concepts of the same type. The concept is defined extensionally, by means of the referents of the mentalesse tokens it produces when the mechanism underlying the ability is properly functioning. Some concepts track more than one thing and are ambiguous, and as I will discuss later some concepts strictly speaking do not track anything and are empty. The particular concept also will have some intension – some set of essential properties of the extension or more typically a set of accidental properties that belong to some members of the extension – used to track the extension. A concept is individuated by its extension, but a conception is individuated both extensionally and intensionally, which is why many types of conceptions correspond to a single type of concept. The intension might be explicitly represented within the person's beliefs, but it does not have to be so. The intension might just be whatever properties of the substance the person is sensitive to and uses in order to identify the substance, in which case it is an empirical question what those properties of the substance are to which the person responds.

The concept's intension does not determine its extension. For instance, I may think of diamonds as hard, colorless, sparkling stones only, but that intension will not exclude cubic zirconia and will not include blue diamonds. What, then, makes all and only diamonds the extension of my concept of diamonds? In this case, the answer is that I have borrowed a term from a public language with public meanings, so mineral experts can fix the extension for me. What about cases in which there are no experts or in which the "experts" are wrong? Suppose that in some earlier time even the "experts" distinguished diamonds from blue diamonds, giving the latter its own name. Would that mean that diamonds are whatever the "experts" say that they are, so there would be no diamonds that were blue? In short, do the *experts'* intensions determine extensions, where an expert is whomever the linguistic community defers to and takes to be an expert? There are shades of Plato's *Euthyphro* here: are diamonds diamonds because the experts say so or do the experts rightly say that something is a diamond because it is a diamond?

If the experts use "diamond" as a name for a natural kind and not merely as an abbreviation for a description (e.g., "whatever is colorless, sparkling, and hard"), then the experts' intension does not necessarily define the extension, and for the same type of reasons that our intension of, say, "Shakespeare" (e.g., as the author of *Hamlet*) does not determine the extension of "Shakespeare": namely, that it is possible that Shakes-

peare did not write *Hamlet* (Kripke 1980). Whether Shakespeare or, say, Bacon wrote *Hamlet* is an historical matter and cannot be settled by semantics (Kripke 1980, 75). Similarly, in the natural sciences, at least, terms such as “diamond” are *supposed* to be rigid designators for natural kinds, although *in fact* some terms such as “phlogiston” are empty and other terms are equivocal. We cannot settle whether all diamonds are necessarily clear by appeal to analytic truths; we must do some scientific investigating (Putnam 1977, 104). This line of thought might seem to imply that I am endorsing a causal theory of reference, but most formulations of such a theory are insufficiently referentially determinate for reasons familiar from Wittgenstein, namely, that no single ostension to christen something as “diamond” (or “gavagai”) will pick out a single substance. We also have to look at how the term is used after its initial ostension:

If I am tracking Fido, I am also tracking the species dog, and also fur and bone. Which of these I am tracking with my mind depends upon which I am learning about or registering information about as I go. And that is determined by which of these substances I identify on other occasions as the one this learning concerns, that is, as being the same substance again. As I dissect my specimen frog in the zoology laboratory, whether I am *conceptually* tracking just the individual Kermit, or tracking frogs, depends on whether I attempt to apply what I have learned from my experience only to later meetings with Kermit or whether to frogs in general. (Millikan 2000, 77)

Even this will not be enough for us to know whether the experts were tracking all diamonds or just colorless diamonds. Insofar as the experts only seemed to be interacting with colorless diamonds when using tokens of “diamond,” we might think that they were “learning about or registering information about” or were identifying as “the same substance again” only colorless diamonds, so that colorless diamonds alone were the extension of “diamond” in their language. Yet there is a strong inclination to think that they were wrong to believe, even in their own language, “All diamonds are colorless.” This correct intuition is founded upon deep biological facts about how we are supposed to form concepts.

Part of the design plan for human cognition is for us to develop certain types of concepts along the lines of natural kinds, where a natural kind is distinguished by “the unlimited richness of the categories [i.e., the amount of unobvious things there are to learn about them], the search for more theory-relevant explanatory properties [than the superficial perceptual properties often used to identify them], the reliance on authority to distinguish exemplars of a category from nonexemplars, the acceptance of abnormal members [e.g., a flightless bird or a blue diamond], and the corrigibility of beliefs about categories” (Markman 1997, 82). There is scientific evidence that humans are genetically programmed to run, when appropriate, inductions on the basis of natural

kinds instead of on the basis of mere perceptual similarity. For instance, four-year olds frequently are prepared to believe that two perceptually dissimilar items are of the same kind; they also are prepared to use the natural kind membership as the basis for further inferences, rather than making the inferences on the basis of misleading perceptual similarities (Markman 1997). It is good for children to be so programmed, because as John Stuart Mill observed, there is much more that we can deduce about something if we know that it is, say, an animal than we could if we merely knew that it was white (Mill 1843, 136). This is not to say that a polar bear, snow, and white paper have nothing interesting in common in virtue of their whiteness, because we can predict that from being white, white things will soil easily, they usually will be seen more easily than black things, etc. (Cangelosi and Parisi 1998; Millikan 2000, 27).⁶ Nonetheless, for many inferences, knowing that something has the property of being white is not particularly helpful, except insofar as whiteness is part of the intension for some natural kind. Colors may help us to make fallible identifications of kinds, but with good reason colors rarely are taken to be essential to the kind – even a polar bear can be painted black. If our imaginary experts once were to have denied that blue diamonds are diamonds, it likely is not because they thought that the color per se was important to the kind but because they thought that in some hidden way colorless diamonds and blue diamonds were importantly different, and that coloring was consequent upon that difference. Compare: we might think that a certain red berry and a certain yellow berry were different kinds of berries because they were different colors, but if we observed that the yellow berries develop into the red ones, we would revise our opinion and think that color differences were not marking a natural kind difference but were just different stages of development within the same kind. A similar type of discovery can be made with the colorless diamonds and blue diamonds, namely, that they are variations of an underlying type – specifically, the blue diamond contains trace amounts of boron.⁷ That children’s thinking naturally develops in the direction of reasoning in terms of natural kinds and that doing so is beneficial implies that this mode of thought is an expression of natural selection and that many types of concepts, including those corresponding to “diamonds,” are supposed to be of natural kinds. The experts’ use of “diamond” was partially wrong, because it failed to correspond to the complete natural kind, and in general the proper function of a substance concept is “to make it possible to utilize substances as these are objectively defined in nature for purposes of gathering and applying information” (Millikan 2000, 50).

At the earliest, immature stages of cognition, most concepts involve perceptual recognition, and this is done without any explicit beliefs. In short, these concepts are know-hows, not know-thats. These recognitional

skills, these know-hows, are the modules discussed earlier. For instance, to have a concept of peaches typically requires learning how to identify peaches directly through the senses, and in the earliest stage of cognitive development this means that a child can make true basic judgments about peaches. These recognitional abilities are not holistic or even molecular. One will have some way of recognizing peaches, but this does not entail any particular beliefs about peaches. A young child may know peaches when she sees them but be unable to say how she recognizes them. In a sense, the child does not *have* beliefs about how to recognize peaches but *embodies* assumptions about what peaches are, assumptions that bridge the gap between the child's stimulations and the child's judgment, and these embodied assumptions are capable of further tuning through experience and training.

The first recognitional abilities seem to precede theoretical concepts and be developed through behavioral reinforcement, but inferential processes also can initiate recognitional abilities. For instance, as a boy Ahab's father might have taught Ahab that whales are gigantic sea creatures. Thus Ahab would have used explicit beliefs in order to identify whales, with the father correcting Ahab as necessary ("That's not a whale; that's a porpoise"), thereby improving Ahab's discriminatory skills. Little Ahab at first might have had to use an inferential process to identify whales, but with time and his father's corrections he could notice other traits of whales, thus refining his discriminatory powers. Some of what he noticed would be represented by explicit beliefs, but some things would simply become part of his whale identification know-how. (I would have a hard time describing well from memory many people whom I know, but I can recognize them instantly.) What began as a labored and less reliable inferential process has been transformed into a more reliable, (usually) effortless, mechanized skill in "seeing" a whale as a whale, and a knowing-that has become a knowing-how.

As I have mentioned, all of these skills are fallible. Young Ahab's skill at identifying whales improves as he gains correction and experience. What makes Ahab's skill be a skill at identifying whales, or be a concept of whales, is not its reliability at identifying whales but the fact that the presence of the skill owes its existence to the fact that it has produced tokens that represent whales in mentalese. The more reliable the skill is, the better, but as it gets more reliable, it still is the same concept, namely, the concept of whales. (As one gets better at riding a bicycle, one is getting better *at riding a bicycle* – one is improving one's skill at doing the same thing.) Reliabilists recognize this point, because they believe that there can be different degrees of reliability at making the same judgment using the same concepts, hence making one's judgment more or less warranted depending upon how it was formed.

So far I have focused on the observational side of concepts, but there is much more to a concept than helping a person recognize it through immediate perception. Much of recognition is done through linguistic reports or through the effects of something not immediately perceived, as when I think of a comet even though I perceive only the crater it made. Also, some concepts are defined in terms of other concepts, so that changing the meaning of these other concepts would change the meaning of the defined concept. Yet even in this case, teleosemantics says that the concept's referent is determined by what its proper function is to map. How and how often it does so map is a different question. Thus the concept's "definition" may not in fact pick out essential properties. Conceptual analysis, ordinary language philosophy, and inferential role semantics reverse the order of priority by making the meaning (intension) of the concept *constitute* what the concept is about, by making the definition *define* reference. By contrast, the teleosemantic approach sees the question of reference as prior and sees meanings (intensions) as fallible, not definitive, means of recognizing the referents. Of course, clarifying definitions and meanings is a legitimate, empirical task in order to understand how and how well the mind cognizes the world. Nonetheless, keeping in mind that from a teleosemantic perspective, the concept is the skill in producing mentalese tokens of a certain type and is not the mentalese tokens themselves, the prior task is to discover what *actual substance* if any is the referent of the tokens in mentalese which the concept has been selected to produce and what is the nature of this substance. To answer these questions requires looking at the receptive side of communication: what actual substance if any does the cognitive system use as the referent of the mentalese tokens, on enough occasions to sustain the joint operations of the production and consumption of these tokens? The "if any" clause is crucial, because empty concepts that produce nonreferring mentalese tokens can be sustained for other reasons. Sometimes the person understands that these tokens are nonreferring, as in literary or folkloric contexts, but sometimes the person believes the tokens to refer. Often these latter concepts are explanatory posits that are defined in terms of other concepts but do not pick out anything, as in the concepts of phlogiston (hypothetical substance that allegedly explains all combustion) and demonic possession (strange behavior allegedly explained by an evil soul that has taken over a person's body). Just as a child could learn to recognize an animal by its tracks, the child could be taught to "recognize" demonic possession by the tracks it seems to leave in the "possessed" person's behavior. The concept is sustained not because it tracks the phenomenon of demonic possession, for there is no such phenomenon, but because it tracks a cluster of observable properties which ostensibly are explained by the possession. Thus such a concept can be useful in

communication (in this case, to pick out behavior) even though the mentalese tokens it produces are nonreferring (or, at least, even though the mentalese tokens do not refer to what people think that they refer to). Indeed, these defective concepts can be modular and involve know-how, so that one just “sees” with no inference that the person is possessed.

A further difference from inferential role semantics is that even though some meanings are “theoretical” in the sense of being defined in terms of other concepts, some concepts are not theoretical in that sense. The ability to recognize a peach as a peach does not have to be theoretical in the sense of being inferential. Of course, in a different sense, all concepts are theoretical, because they might “be profoundly inappropriate to reality” (Churchland 1979, 41). From this discussion it should be clear why I think that in a certain sense all concepts are theory-laden but in another sense many are not.⁸

Notes

¹ A complication which I will discuss later is that beliefs about appearances fall into two different classes: first-order and second-order appearance beliefs. Yet their epistemological function is the same: to serve as basic beliefs that can be used to make inferences about how things are.

² As I will discuss later, this does not mean that other beliefs never affect the mechanism.

³ This picture of how the capacity for thoughts about appearances develops is indebted to Sellars’ story about the Rylean ancestors who began as behaviorists and later posited that there are thoughts and inner impressions (2000, sec. 48 ff., pp. 258 ff.). Doing so enabled them to explain why, for instance, someone will report that something is red even if the thing is not red or even if there is nothing at all. This is not to deny that these Rylean ancestors could not learn to notice or introspect impressions after their existence had been posited – in such a case, “*What began as a language with a purely theoretical use has gained a reporting role*” (Sellars 2000, sec. 59, p. 269). The boundary between what is observed and what is deduced on the basis of observation is dynamic. See also Millikan (2000, 86-7) on the potential variability of what is directly perceived despite constancy of stimulation and stimulus.

⁴ Think of concepts such as of salt and of pepper or concepts associated as cause and effect. There are other types of associations of opposites that would not have the same psychological effect (e.g., of straight/curved, white/black, tall/short, etc.).

⁵ “Holistic properties are properties such that, if anything has them, then *lots* of other things must have them too.” (Fodor and Lepore 1992, 2)

⁶ It might be thought that there must be some interesting but hidden properties of all white things; we still do not fully understand the interaction of our visual system with light and of light with pigmentation, so we do not really understand what it is for something to be white. Yet, this reductive line of thought is

unpromising, because it seems unlikely that white things share any distinctive property aside from their power to make us see white (or to reflect white light). Being white does not supervene on any univocal type of state of affairs – it is a “disjunctive property.”

⁷ See the American Museum of Natural History’s “The Nature of Diamonds” at www.amnh.org/exhibitions/diamonds/color.html and *Encyclopedia Smithsonian*, s.v. “Hope Diamond,” www.si.edu/resource/faq/nmnh/hope.htm.

⁸ Special thanks go to Ruth Millikan for her comments on an earlier draft of this paper and to the participants of the Conference on Coherentism in Epistemology for their questions and comments.

Bibliography

- Cangelosi, Angelo and Domenico Parisi. 1998. Concepts in Artificial Organisms. *Behavioral and Brain Sciences* 21, no. 1: 68-9.
- Churchland, Paul M. 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge Studies in Philosophy. Cambridge: Cambridge University Press.
- Fodor, Jerry and Ernest Lepore. 1992. *Holism: A Shopper’s Guide*. Oxford: Blackwell.
- Harman, Gilbert. 1997. Positive Versus Negative Undermining in Belief Revision. In *Naturalizing Epistemology*, 2nd ed., ed. Hilary Kornblith (Cambridge: A Bradford Book, MIT Press). First published in *Change in View: Principles of Reasoning* (Cambridge: MIT Press, 1986).
- Kripke, Saul A. 1980. *Naming and Necessity*. Cambridge: Harvard University Press.
- Markman, Ellen. 1997. Natural kinds. In *Naturalizing Epistemology*, 2nd ed., ed. Hilary Kornblith (Cambridge, MA: MIT Press, A Bradford Book). Reprinted from *Categorization and Naming in Children* (Cambridge: MIT Press, 1989).
- Mill, John Stuart. 1843. *A System of Logic, Ratiocinative and Inductive*. London: Longmans.
- Millikan, Ruth Garrett. 2000. *On Clear and Confused Ideas: An Essay About Substance Concepts*. Cambridge Studies in Philosophy. Cambridge: Cambridge University Press.
- Sellars, Wilfred. 2000. Empiricism and the Philosophy of Mind. In Willem A. de Vries and Timm Triplett, *Knowledge, Mind, and the Given: Rereading Wilfred Sellars’s “Empiricism and the Philosophy of Mind,” Including the Complete Text of Sellars’s Essay* (Indianapolis: Hackett). Originally published in H. Fiegl and M. Scriven, ed., *Minnesota Studies in the Philosophy of Science*, vol. 1 (Minneapolis: University of Minnesota Press, 1956).