The Centrality of Sellars's Two-Ply Account of Observation to the Arguments of "Empiricism and the Philosophy of Mind"

"Empiricism and the Philosophy of Mind" is one of the great works of twentieth century philosophy. It is rich, deep, and revolutionary in its consequences. It cannot, however, be ranked among the most *perspicuous* of philosophical writings. Although it is fairly easy to discern its general tenor and tendency, the convoluted and digressive order of exposition pursued in the essay has obscured for many readers the exact outlines of such a fundamental concept as *givenness*—with the result that few could at the end of their reading accurately trace its boundaries and say what all its species have in common, being obliged instead to content themselves with being able to recognize some of its exemplary instances. Again, I think that partly for this reason, readers of *EPM* seldom realize just how radical is its critique of empiricism—just how much of traditional empiricist ways of thinking must be rejected if Sellars' arguments are accepted. And if the full extent of the work's conclusions is hard to appreciate, all the more difficult is it to follow its argumentative path through all its turnings. In what follows my aim is to lay out one basic idea of Sellars's, which I see as underlying three of the most important arguments he deploys along the way to his conclusions. My concern here will not be in how those arguments contribute to his overall enterprise, but rather in how they are rooted in a common thought. Sellars does not make this basic idea as explicit as one would like, and does not stop along the way to observe how each of the three individual arguments depends on it. But if I am right, we will understand the essay better by being able to identify and individuate this thread in the tapestry.

The master idea I want to start with is Sellars's understanding of observational capacities: the ability to make noninferential reports of, or to form perceptual judgments concerning, perceptible facts. My claim is that he treats them as the product of two distinguishable sorts of abilities: the capacity reliably to discriminate behaviorally between different sorts of stimuli, and the capacity to take up a position in the game of giving and asking for reasons. The three central strategic moves in the essay I will seek to understand in terms of that two factor approach to observation are: first, the way he dissolves a particular cartesian temptation by offering a novel account of the expressive function of 'looks' talk; second, his rationalist account of the acquisition of empirical concepts; and third, his account of how theoretical concepts can come to have observational uses.

I. Sellars's Two-Ply Account of Observation

If we strip empiricism down to its core, we might identify it with the insight that knowledge of the empirical world depends essentially on the capacity of knowing organisms to respond differentially to distinct environing stimuli. I'll call this claim '*basic'*, or 'stripped down' empiricism; it could equally well be called the *trivial* thesis of empiricism.[1] Surely no rationalist or idealist has ever denied *this* claim. While differential responsiveness is obviously a necessary condition for empirical knowledge, it is clearly nothing like a sufficient condition. A chunk of iron responds differentially to stimuli, for instance by rusting in some environments, and not in others. To that extent, it can be construed as *classifying* its environments, taking or treating them as being of one of two kinds. In the same way, as Hegel says, an animal takes

[1] I would call it 'minimal empiricism', except that John McDowell, in the Introduction to the paperback edition of *Mind and World* (Harvard University Press, 1996) has adopted that term for a *much* more committal thesis.

something as food by "falling to without further ado and eating it up."[2] But this sort of classificatory taking something *as* something should not yet be classed as a *cognitive* matter, on pain of losing sight of the fundamental ways in which genuine observationally acquired knowledge differs from what is exhibited by merely irritable devices such as thermostats and land mines.

A parrot could be trained to respond to the visible presence of red things by uttering the noise "That's red." We might suppose that it is disposed to produce this performance under just the same circumstances in which a genuine observer and reporter of red things is disposed to produce a physically similar performance. There is an important respect in which the parrot and the observer are alike. We could call what they share a *reliable differential responsive disposition* (which I'll sometimes shorten to '*RDRD*'). RDRDs are the first element in Sellars's two-ply account of observational knowledge. At least in the basic case, they are characterizable in a naturalistic, physicalistic vocabulary.[3] The concept of an RDRD is meant to capture the capacity we genuine knowers share with artifacts and merely sentient creatures such as parrots that the basic thesis of empiricism insists is a necessary condition of empirical knowledge.

The second element of Sellars's two-ply account of observational knowledge is meant to distinguish possessors of genuine observational belief and knowledge from merely reliable differential responders. What is the crucial difference between the red-discriminating parrot and the genuine observer of red things? It is the difference between *sentience* and *sapience*. For

[2] *Phenomenology*, paragraph 109, in the numeration of A.V. Miller's translation (Oxford University Press, 1979).
[3] They would not be so characterizable in cases where the response is specified in, say, normative or semantic vocabulary—for instance, as *correctly* using the word 'red', or as applying the *concept* red.

Sellars's purposes in EPM, the difference between merely differentially responding artifacts and genuinely sentient organisms does not make an essential cognitive or epistemological difference. *All* we need pay attention to in them is their exercising of reliable differential responsive dispositions. But he is very concerned with what distinguishes both of these sorts of things from genuine observers. His thought is that the difference that makes a difference is that candidates for observational knowledge don't just have reliable dispositions to respond differentially to stimuli by *making noises*, but have reliable dispositions to respond differentially to those stimuli by *applying concepts*. The genuine observer responds to visible red things by coming to believe, claiming, or reporting *that* there is something red. Sapient awareness differs from awareness in the sense of mere differential responsiveness (the sort exhibited by any organism or device that can for instance be said in the full sense to be capable of avoiding obstacles) in that the sapient being responsively classifies the stimuli as falling under concepts, as being of some conceptually articulated kind.

It is obvious that everything turns on how one goes on to understand concept application or the conceptual articulation of responses. For Sellars, it is a linguistic affair: grasping a concept is mastering the use of a word. Then we must ask what makes something a use of a word, in the sense relevant to the application of concepts. Sellars's answer is that for the response reliably differentially elicited by the visible presence of a perceptible state of affairs to count as the application of a concept, for it to be properly characterized as a reporting or coming to believe *that* such-and-such is the case, is for it to be the making of a certain kind of move or the taking up of a certain kind of position in a game of giving and asking for reasons. It must be committing oneself to a content that can both serve as and stand in need of *reasons*, that is, that

can play the role both of premise and of conclusion in *inferences*. The observer's response is conceptually contentful just insofar as it occupies a node in a web of inferential relations.

What the parrot lacks is a *conceptual understanding* of its response. That is why it is just making noise. Its response means nothing to the parrot—though it may means something to us, who *can* make inferences from it, in the way we do from changes in the states of measuring instruments. The parrot does not treat red as entailing colored, as entailed by scarlet, as incompatible with green, and so on. And because it does not, uttering the noise 'red' is not, for the parrot, the adopting of a stance that can serve as a reason committing or entitling it to adopt other stances, and potentially in need of reasons that might be supplied by still further such stances. By contrast, the observer's utterance of 'That's red,' is making a move, adopting a position, in a game of giving and asking for reasons. And the observer's grasp of the conceptual content expressed by her utterance consists in her practical mastery of its significance in that game: her knowing (in the sense of being able practically to discriminate, a kind of knowing *how*) what follows from her claim and what it follows from, what would be evidence for it and what is incompatible with it.

Although Sellars does not carefully distinguish them, two different strands can be discerned within this second element of his account. First is the idea that for performances (whether noninferentially elicited responses or not) to count as *claims*, and so as expressions of *beliefs* or *judgments*, as candidates for *knowledge*, they must be in what he calls "the dimension

of endorsement."[4] This is to say that they must have a certain sort of pragmatic significance or force: they must express the endorsement of some content by the candidate knower. They must be the adoption of a certain kind of normative stance: the undertaking of a *commitment*. Second, that the commitment is a *cognitive* commitment, the endorsement of a *conceptual content*, is to be understood in terms of its *inferential* articulation, its place in the "space of reasons," its being a move in the "game of giving and asking for reasons."[5] This is to say at least that in making a claim one commits oneself to its suitability as a premise from which conclusions can be drawn, a commitment whose entitlement is always at least potentially liable to demands for vindication by the exhibition of other claims that can serve as reasons for it.

This two factor account of perceptual judgments (claims to observational knowledge) is a version of a broadly kantian strategy: insisting on the collaboration of capacities characterizable in terms of receptivity and spontaneity. It is a pragmatic version, since it is couched in terms of know *how*: practical abilities to respond differentially to nonlinguistic stimuli, and to distinguish in practice what inferentially follows from or serves as a reason for what. The residual empiricism of the approach consists in its insistence on the need for the exercise of some of our conceptual capacities to be the exercise of RDRDs. Its residual rationalism consists in its insistence that the responses in question have cognitive significance, count as applications of concepts, only in virtue of their role in reasoning. What otherwise would appear as language-entry moves, without language-language moves, are blind. What otherwise would appear as language-language moves without language-entry moves, are empty. (I say "what otherwise

[4] Sellars discussion begins at [EPM16]. All references are to section numbers of Sellars' "Empiricism and the Philosophy of Mind", reprinted with an Introduction by Richard Rorty and a Study Guide by Robert Brandom (Harvard University Press, 1997).

[5] See for instance [EPM36].

would appear" as moves because such blind or empty moves do not for Sellars qualify as moves in a *language* game at all.)[6]

It follows from this two-pronged approach that we must be careful in characterizing perceptual judgments or reports of observations as 'noninferential'. They are noninferential in the sense that the particular acts or tokenings are noninferentially *elicited*. They are not the products of a process of inference, arising rather by the exercise of reliable capacities to noninferentially respond differentially to various sorts of perceptible states of affairs by applying concepts. But *no* beliefs, judgments, reports, or claims—in general, no applications of concepts—are noninferential in the sense that their content can be understood apart from their role in reasoning as potential premises and conclusions of inferences. Any response that does not at least potentially have an inferential significance—which cannot, for instance serve as a premise in reasoning to further conclusions—is cognitively idle: a wheel on which nothing else turns.

This rationalist claim has radical consequences. It means that there can be no language consisting only of noninferential reports, no system of concepts whose *only* use is in making perceptual judgments. Noninferential reports do not form an autonomous stratum of language: a game one could play though one played no other. For that they are *reports* or *claims*, expressions of *beliefs* or *judgments*, that they are applications of *concepts* at all, consists in their availability to serve as premises and conclusions of inferences. And this is so no matter what the

[6] The idiom of "language-language" moves and "language-entry" moves is drawn from Sellars' "Some Reflections on Language Games",in *Science, Perception, and Reality* (London: Routledge, Kegan Paul, 1963). Reprinted in *In the Space of Reasons*.

subject matter of the reports might be—even if what is reported, that of which one is noninferentially aware, is one's own current mental states. Awareness that reaches beyond mere differential responsiveness—that is, awareness in the sense that bears on *cognition*—is an essentially inferentially articulated affair.

So observational concepts, ones that have (at least some) noninferential circumstances of appropriate application, can be thought of as *inference laden*. It does not follow, by the way, that they are for Sellars for that reason also *theory laden*. For, as will appear below, Sellars understands theoretical concepts as those that have only inferential circumstances of appropriate application—so that noncompound claims in which they occur essentially are ones that one can only become entitled to as the result of an inference. His rationalist rendering of the notion of conceptual contentfulness in terms of role in reasoning only commits Sellars to the claim that for any concept to have noninferential uses, it must have inferential ones as well. He is prepared to countenance the possibility of an autonomous language game in which every concept has noninferential, as well as inferential uses. Such a language game would be devoid of theoretical terms.

II. 'Looks' Talk and Sellars's Diagnosis of the Cartesian Hypostatization of Appearances

One of the central arguments of EPM applies this two-legged understanding of the use of observational concepts to the traditional understanding of claims about how things *look* as reports of *appearances*. The question he addresses can be variously put. In one form it is the question of whether <u>looks-red</u> come before <u>is-red</u> conceptually (and so in the order of explanation)? Put in a form more congenial and comprehensible to a pragmatist—that is, in a

form that concerns our abilities to *do* something—this becomes the question of whether the latter can be defined in terms of the former in such a way that one could learn how to use the defining concept (looking-φ) first, and only afterwards, by means of the definition, learn how to use the defined concept (is-φ)? Since Sellars understands grasp of a concept in terms of mastery of the use of a word, this then becomes a question about the relation between practices of using "looks-φ" talk and the practices of using "is-φ" talk. This is a relatively clear way of asking about an issue that goes to the heart of the cartesian project of defining the ontological realm of the mental in terms of the epistemic privileged accesss in the sense of incorrigibility of mental occurrences.

Descartes was struck by the fact that the appearance/reality distinction seems not to apply to appearances. While I may be mistaken about whether something *is* red (or whether the tower, in the distance, *is* square), I cannot in the same way be mistaken about whether it *looks* red to me now.[7] While I may legitimately be challenged by a doubter: "Perhaps the item is not *really* red; perhaps it only *seems* red," there is no room for the further doubt, "Perhaps the item does not even *seem* red; perhaps it only *seems* to seem red." If it seems to seem red, then it really does seem red. The *looks*, *seems*, or *appears* operators collapse if we try to iterate them. A contrast between appearance and reality is marked by the distinction between looks-φ and φ for ordinary (reality-indicating) predicates 'φ'. But no corresponding contrast is marked by the distinction between looks-to-look-φ and looks-φ. Appearances are reified by Descartes as things that really

[7] I might be mistaken about whether *red* is what it looks, that is, whether the property expressed by the word 'red' is the one it looks to have. But that, the thought goes, is another matter. I cannot be mistaken that it looks that way, like *that*, where this latter phrase is understood as having a noncomparative use. It *looks-red*, a distinctive phenomenal property, which we may inconveniently only happen to be able to pick out by its association with a word for a real-world property.

are just however they appear. He inferred that we do not know them mediately, by means of representings that introduce the possibility of *mis*-representing (a distinction between how they really are and how they merely appear, i.e. are represented as being). Rather, we know them *immediately*—simply by having them. Thus appearings—thought of as a realm of entities *reported* on by noninferentially elicited claims about how things *look* (for the visual case), or more generally *seem*, or *appear*—show up as having the ideal qualifications for epistemologically secure foundations of knowledge: we cannot make mistakes about them. Just *having* an appearance ("being appeared-to φ-ly", in one of the variations Sellars discusses) counts as *knowing* something: not that something is φ, to be sure, but at least that something *looks-, seems-,* or *appears*-φ. The possibility accordingly arises of reconstructing our knowledge by starting out only with knowledge of this sort—knowledge of how things look, seem, or appear—and building up in some way to our knowledge (if any) of how things really are (outside the realm of appearance).

This project requires that concepts of the form <u>looks-φ</u> be intelligible in principle in advance of grasping the corresponding concepts <u>φ</u> (or <u>is-φ</u>). Sellars argues that Descartes got things backwards. 'Looks' talk does not form an autonomous stratum of the language—it is not a language-game one could play though one played no other. One must already be able to use 'is-φ' talk in order to master 'looks-φ' talk, which turns out to be parasitic on it. In this precise practical sense, <u>is-φ</u> is *conceptually* (Sellars often says 'logically') *prior* to <u>looks-φ</u>.

His argument takes the form of an account of how 'looks' talk can arise piggy-backed on 'is' talk. In *EPM* Sellars does not try to support the strong modal claim that the various practices

must be related in this way. He thinks that his alternative account of the relation between these idioms is so persuasive that we will no longer be tempted by the cartesian picture. It is an interesting question, which I will not pursue here, whether his story can be turned into an more compelling argument for the stronger claim he wants to make. What he offers us is the parable of John in the tie shop.

At the first stage, John has mastered the noninferential use of terms such as 'green' and 'blue'. So he can, typically, reliably respond to green things by applying the concept green, to blue things by applying the concept blue, and so on. To say that his responsive dispositions are reliable is to say that he usually turns out to be right—so the inference from his being disposed to call something 'green' or 'blue' to its being green or blue is a generally good (though not infallible) one.

At the next stage, electric lights are installed in the shop, and John discovers that they make him prey to certain sorts of systematic errors. Often, when under the electric lights inside his shop he observes something to be green, it turns out in fact—when he and others examine it outside in daylight—to be blue. Here it is obviously important that John have access to some ways of entitling himself to the claim that something is blue, besides the term he is initially disposed to apply to it. This can include his dispositions to respond to it outside the shop, together with his beliefs about the circumstances in which ties do and do not change color, the assessments of others, and the fact that the proper use of color terms was originally keyed to daylight assessments. At this point, John becomes cautious. When viewing under the nonstandard conditions of electric lighting, he does not indulge his otherwise reliable disposition to respond

to some visible ties by calling them green. Instead he says something like: "I'm disposed to call this green, and if I didn't know that under these circumstances I'm not a reliable discriminator of green things, I would give in to that temptation and call it green."

At the final stage, John learns under these circumstance to substitute the expression "It *looks* green," for this long expression of temptation withstood. Using the expression "looks-φ" is doing two things: first, it is evincing the same usually reliable differential responsive disposition that in other circumstances results in the claim that something *is* (. But second, it is *withholding* the endorsement of the claim that something is green. In other words, it is doing something that agrees with an ordinary non-inferential report of green things on the first component of Sellars's two-ply account of observation reports—sharing an RDRD—but disagrees with it on the second component, withholding endorsement instead of undertaking the commitment.

The idea is that where collateral beliefs indicate that systematic error is likely, the subject learns not to make the report '*x* is φ', to which his previously inculcated responsive dispositions incline him, but to make a new kind of claim: '*x looks* (or seems) φ'. The cartesian temptation is to take this as a new kind of report. This report then is naturally thought of as reporting a minimal, noninferentially ascertainable, foundationally basic item, an appearing, about which each subject is incorrigible. Sellars's claim is that it is a mistake to treat these as reports at all—since they *evince* a disposition to call something φ, but do not do so. They do not even *report* the presence of the disposition—that is, they are not ways of *saying* that one has that disposition.

This analysis of what one is doing in using 'looks' explains the incorrigibility of 'looks' talk. One can be wrong about whether something *is* green because the claim one endorses, the commitment one undertakes, may turn out to be incorrect. For instance, its inferential consequences may be incompatible with other facts one is or comes to be in a position to know independently. But in saying that something *looks* green, one is not endorsing a claim, but *withholding* endorsement from one. Such a reporter is merely evincing a disposition to do something that for other reasons (e.g. suspicion that the circumstances of observation lead to systematic error) he is unwilling to do—namely, endorse a claim. Such a reporter cannot be wrong, because he has held back from making a commitment. This is why the *looks*, *seems*, and *appears* operators do not iterate. Their function is to express the withholding of endorsement from the sentence that appears within the scope of the operator. There is no sensible contrast between 'looks-to-look φ' and 'looks-φ', of the sort there is between 'looks-φ' and '(is-φ (' because the first 'looks' has already withheld endorsement from the only content in the vicinity to which one might be committed (to something's being φ). There is no further withholding work for the second 'looks' to do. There is nothing left to take back. Since asserting 'X looks φ' is not undertaking a propositionally contentful commitment—but only expressing an overrideable disposition to do so—there is no issue as to whether or not that commitment (which one?) is correct.

Sellars accordingly explains the incorrigibility of appearance-claims, which had so impressed Descartes. He does so in terms of the practices of using words, which are what grasp of the relevant appearance concepts must amount to, according to his methodological linguistic pragmatism. But once we have seen the source and nature of this incorrigibility—in down-to-

earth, practical, resolutely nonmetaphysical terms—we see also why it is precisely unsuited to use as an epistemological foundation for the rest of our (risky, corrigible) empirical knowledge. For, first, the incorrigibility of claims about how things merely *look* simply reflects their emptiness: the fact that they are not really claims at all. And second, the same story shows us that 'looks' talk is not an autonomous language game—one that could be played though one played no other. It is entirely parasitic on the practice of making risky empirical reports of how things actually are. Thus Descartes seized on a genuine phenomenon—the incorrigibility of claims about appearances, reflecting the non-iterability of operators like *looks, seems,* and *appears*—but misunderstood its nature, and so mistakenly thought it available to play an epistemologically foundational role for which it is in no way suited.

III. Two Confirmations of the Analysis of 'Looks' Talk in Terms of the Two-Ply Account of Observation

Sellars finds that the analysis of 'looks' talk in terms of the two pronged account of perceptual judgments is confirmed by its capacity to explain features of appearance-talk that are mysterious on the contrasting cartesian approach.

i) The apple over there is red.

ii) The apple over there looks red.

iii) It looks as though there were a red apple over there.

Utterances of these sentences can express the same responsive disposition to report the presence of a red apple, but they endorse (take responsibility for the inferential consequences of) different parts of that claim. (i) endorses both the existence of the apple, and its quality of redness. (ii) endorses only the existence of the apple. The 'looks' locution explicitly cancels the qualitative

commitment or endorsement. (iii) explicitly cancels both the existential and the qualitative endorsements. Thus, if someone claims that there is in fact no apple over there, he is asserting something incompatible with (i) and (ii), but not with (iii). If he denies that there is anything red over there, he asserts something incompatible with (i), but not with (ii) or (iii). Sellars's account of the practice of using 'looks', in terms of the withholding of endorsement when one suspects systematic error in one's responsive dispositions, can account for the difference in scope of endorsement that (i)-(iii) exhibit. But how could that difference be accounted for by an approach that understands 'looks' talk as reporting a distinctive kind of particular, about which we are incorrigible?

Sellars finds a further confirmation of his account of 'looks' talk—and so of the two factor account of observational capacities that animates it—in its capacity to explain the possibility of reporting a merely *generic* (more accurately, merely determinable) look. Thus it is possible for an apple to look red, without its looking any specific shade of red (crimson, scarlet, etc.). It is possible for a plane figure to look many-sided without there being some particular number of sides (say 119) which it looks to have. But if 'looks' statements are to be understood as reports of the presence before the eye of the mind of a particular which *is* (, how can this possibility be understood? Particulars are completely determinate. A horse has a particular number of hairs, though as Sellars points out, it can *look* to have merely 'a lot' of them. It is a particular shade of brown (or several shades), even though it may look only darkly colored. So how are such generic, merely determinable, looks possible? Sellars's account is in terms of scope of endorsement. One says that the plane figure looks 'many-sided' instead of '119-sided' just in case one is willing only to endorse (be held responsible for justifying) the more general claim. This is

a matter of how far one is willing to trust one's responsive dispositions, a matter of the epistemic credence one feels they deserve or are able to sustain. Particulars, even if they are sense contents, cannot be colored without being some determinate color and shade. How then can the sense datum theorist—who wants to say that when something *looks* φ to S, something in S *is* φ—account for the fact that something can look colored without looking to be any particular color, or look red without looking to be any particular shade of red? So Sellars's account of 'looks' talk in terms of endorsement can account for two aspects of that kind of discourse that no theory that invokes a given can explain: the scope distinctions between qualitative and existential lookings, and the possibility of merely generic or determinable lookings.

IV. A Rationalist Account of the Acquisition of Empirical Concepts

It is characteristic of empiricism as Sellars understands (and rejects) it, that it countenances a notion of awareness or experience meeting two conditions. First, it goes beyond mere differential responsiveness in having some sort of cognitive *content*—that is, content of the sort that under favorable circumstances amounts to knowledge. This is the idea of a notion of awareness or experience *of* a red triangle in one's visual field that can at the same time be (or be one's evidence for) knowledge *that* there is a red triangle in one's visual field. Second, this sort of awareness is *preconceptual*: the capacity to be aware in this sense or have experiences of this sort is prior to and independent of the possession of or capacity to apply concepts. The idea of a kind of awareness with these two features is what Sellars calls the "Myth of the Given."

Whatever difficulties there may be with such a conception—most notably the incoherences Sellars rehearses in the opening sections of EPM—it does provide the basis for a story about concept acquisition. Concepts are understood as acquired by a process of *abstraction*, whose raw materials are provided by exercises of the primitive capacity for immediate, preconceptual awareness.[8] One may—and Sellars does—raise questions about whether it is possible to elaborate this story in a coherent fashion. But one ought also to ask the corresponding question to the empricists' rationalist opponents. Rationalists like Sellars claim that *all* awareness is a conceptual affair. Being aware of something, in any sense that goes beyond mere responsiveness in its potential cognitive significance—paradigmatically in its capacity to serve as *evidence*—is bringing it under a concept. Sense experience cannot be the basis for the acquisition of concepts, since it presupposes the capacity to apply concepts. So how *do* knowers acquire concepts? At this point in the dialectic, classical rationalists such as Leibniz threw up their hands and invoked innate ideas—denying that at least the most basic and general concepts *were* acquired at all. Sellars owes either a defense of innatism, or an alternative account of concept acquisition.

Sellars rejects innatism. Grasp of a concept is mastery of the use of a word, so concepts are acquired in the process of learning a language. But if we don't acquire the concept <u>green</u> by noticing green things, since we must already have the concept in order to notice green things as

[8] It is tempting to think that on this line concepts are related to the contents of preconceptual experiences as universals to particulars. But as Sellars points out, the empiricists in fact took as primitive the capacity to be aware of already of *repeatables*, such as redness and squareness. This might suggest that the relation is better understood as one of genus to species. But <u>scarlet</u> is not strictly a *species* of the genus <u>red</u>, since there need be no way to specify the relevant differentiae without mentioning the species. (Compare the relation between the phenomenal property of redness and that of being colored.) So the relation between immediately experienceable contents and the concepts under which they are classified is better understood as that of *determinate* repeatable to *determinables* under which it falls.

such (by applying the concept to them), how is it possible for us to learn the use of the word 'green', and hence acquire the concept? We each start by learning the corresponding RDRDs: being trained to respond to visibly green things by uttering what is still for the novice just the *noise* 'green'. This much, the parrot can share. Besides these language-entry moves, the language learner must also master the inferential moves in the vicinity of 'green': that the move to 'colored' is OK, and the move to 'red' is not, and so on. Training in these basic language-language moves consists in acquiring more RDRDs, only now the stimuli, as well as the responses, are utterances.

If a two year old wobbles into the living room and utters the sentence "The house is on fire," we will not generally take him to have claimed or expressed the belief that the house is on fire. He does not know what he is saying—in the sense that he does not yet know what he would be committing herself to by that claim, or what would be evidence for it or against it. If a five year old child utters the same sentence, though, we may well take the utterance to have the significance of a claim, the expression of a belief. We take it to be the adoption of a stance in the dimension of endorsement, to be the undertaking of a commitment, by *holding* the child responsible for her claim: asking for her evidence, asking her what she thinks we should to about it, and so on. For it is now presumed that she can tell what she is committing herself to, and what would entitle her to that commitment, and so knows what she is saying, what claim she is endorsing, what belief she is expressing. When the child masters enough of the inferential moves in the vicinity of a responsively elicited utterance of "That is red," she is taken to have endorsed a claim, and so to have applied a concept.

On the inferential account of distinctively conceptual articulation, grasping a concept requires mastering the inferential connections between the appropriate use of some words and the appropriate use of others. So on this account there is no such thing as grasping just one concept: grasping *any* concept requires grasping *many* concepts. Light dawns slowly over the whole.

How good must one be at discriminating the appropriate antecedents and consequents of using a word in order to count as grasping the concept it expresses? Sellars does not explicitly address this question in EPM, but I think his view is that whether or not one's utterance has the significance of endorsing a claim, and so of applying a concept, is a question of how it is treated by the other members of the linguistic community. The normative status of committing oneself—taking up a position in the dimension of endorsement—is a social status. One must be good enough at anticipating and fulfilling one's responsibilities in order to be *held* responsible, and so for one's remarks to be accorded authority, in the sense of being treated as providing suitable premises for inferences by others. How much is enough is not a metaphysical matter of recognizing the crossing of some antecedently specifiable boundary, but a social matter of deciding when to recognize a performance as authoritative and hold the performer responsible. It is a question that belongs in a box with when writing one's name at the bottom of a piece of paper counts as committing oneself to pay the bank a certain sum of money every month for thirty years. Some seventeen year olds may actually understand what they would be committing themselves to better than some twenty two year olds. But the community is not therefore making a metaphysical mistake in treating the latter but not the former as able genuinely to commit themselves.

Sellars account of concept acquisition starts with reliable differential responsive dispositions to respond to environing stimuli by uttering sentences. What is then required is that one's utterance come to have the significance of making a move in the game of giving and asking for reasons. That requires two elements: the practical inferential know-how required to find one's way about in the inferential network connecting different sentences, and the social acknowledgment of that know-how as sufficient for one's performances to have the significance in the linguistic community of commitments to or endorsements of the inferentially articulated claims expressed by those sentences. This story is structured and motivated by Sellars's two-pronged account of observation reports, as non-inferentially elicited endorsements of inferentially articulated claims.

V. Giving Theoretical Concepts an Observational Use

As a final example of the work Sellars calls on his two pronged analysis of observational capacities to do in "Empiricism and the Philosophy of Mind", we might consider his account of how theoretical concepts can acquire an observational use. His reason for addressing the issue is that he wants to make intelligible the idea that some sorts of paradigmatic mental occurrences—thoughts and sense impressions—might first become available to us purely theoretically, and only later come to be observable by us. For showing that such a development in our capacities *is* intelligible provides a means of confounding the cartesian idea of immediate (that is, noninferential) observability as *essential* to the very idea of mental occurrences. But my concern here is with the general point, rather than this particular application of it.

The first point to realize is that, as I mentioned above, according to Sellars's view, the distinction between theoretical objects and observable objects is *methodological*, rather than *ontological*. That is, theoretical and observable objects are not different kinds of thing. They differ only in how we come to know about them. Theoretical objects are ones of which we can only have *inferential* knowledge, while observable objects can also be known noninferentially. Theoretical concepts are ones we can only be entitled to apply as the conclusions of inferences, while concepts of observables also have noninferential uses. But the line between things to which we have only inferential cognitive access and things to which we also have noninferential cognitive access can shift with time, for instance as new instruments are developed. Thus when first postulated to explain perturbations in the orbit of Neptune, Pluto was a purely theoretical object; the only claims we could make about it were the conclusions of inferences. But the development of more powerful telescopes eventually made it accessible to observation, and so a subject of noninferential reports. Pluto did not then undergo an ontological change. All that changed was its relation to us.[9]

It might be objected to this view that when the issue of the ontological status of theoretical entities is raised, they are not considered merely as objects in principle like any others save that they happen at the moment to be beyond our powers of observation. They are thought of as *unobservable* in a much stronger sense: permanently and in principle inaccessible to observation. But Sellars denies that anything is unobservable in this sense. To be observable is

[9] Notice that this realism about theoretical entities does not entail scientific realism in the sense that privileges science over other sorts of cognitive activity, although Sellars usually discusses the two sorts of claims together.

just to be noninferentially reportable. Noninferential reportability requires only that there are circumstances in which reporters can apply the concepts in question (the dimension of inferentially articulated endorsement) by exercising reliable differential dispositions to respond to the objects in question (the causal dimension), and know that they are doing so. In this sense, physicists with the right training can *noninferentially* report the presence of mu mesons in cloud chambers. In this sense of 'observation', nothing real is in principle beyond the reach of observation. (Indeed, in Sellars's sense, one who mastered reliable differential responsive dispositions noninferentially to apply *normative* vocabulary would be directly observing normative facts. It is in this sense that we might be said to be able to *hear*, not just the noises someone else makes, but their *words*, and indeed, *what they are saying*—their *meanings*.) It is an empirical question what circumstances we can come reliably to respond to differentially. The development of each new sort of measuring instrument potentially expands the realm of the here-and-now observable.

Once one sees that observation is not based on some primitive sort of preconceptual awareness, the fact that some observation reports are riskier than others and that when challenged we sometimes retreat to safer ones from which the originals can be inferred will not tempt one to think that the original reports were in fact the products of inference from those basic or minimal observations. The physicist, if challenged to back up his report of a mu-meson may indeed justify his claim by citing the distinctively hooked vapor trail in the cloud chamber. This is something else observable, from which the presence of the mu meson can, in the right circumstances, be inferred. But to say that is not to say that the original report was the product of an inference after all. It was the exercise of a reliable differential responsive disposition

keyed to a whole chain of reliably covarying events, which includes mu mesons, hooked vapor trails, and retinal images. What makes it a report of mu mesons, and not of hooked vapor trails or retinal images is the inferential role of the concept the physicist noninferentially applies. (It is a consequence of something's being a mu meson, for instance, that it is *much* smaller than a finger, which does *not* follow from something's being a hooked vapor trail.) If *mu meson* is the concept the physicist applies noninferentially, then if he is sufficiently reliable, when correct, that is what he *sees*. His retreat, when a question is raised, to a report of a hooked vapor trail, whose presence provides good inferential reason for the original, noninferentially elicited claim, is a retreat to a report that is safer in the sense that he is a *more* reliable reporter of hooked vapor trails than of mu mesons, and that it takes less training to be able reliably to report vapor trails of a certain shape, so that is a skill shared more widely. But the fact that an inferential justification can be offered, and that the demand for one may be in order, no more undermines the status of the original report as noninferentially elicited (as genuinely an observation) than does the corresponding fact that I may under various circumstances be obliged to back up my report of something as red by invoking my reliability as a reporter of red things in these circumstances—from which, together with my disposition to call it red, the claim originally endorsed noninferentially may be inferred.

Thus one can start with grasp of a concept that consists entirely in mastery of its use as a premise and conclusion in inferences—that is, as a purely theoretical concept—and by the addition of suitable RDRDs come to be able to use them observationally, perhaps in observations whose standard conditions include not only such items as good light (as in the tie shop case) but also the presence of various sorts of instruments. This argument once again appeals to and

depends upon Sellars's understanding of observational capacities as the product of reliable noninferential responsive dispositions and mastery of inferential norms.

VI. Conclusion: On the Relation Between the Two Components

Sellars's primary explanatory target in "Empiricism and the Philosophy of Mind" is our knowledge of the current contents of our own minds. He wants to rethink our understanding of the way in which we experience or are aware of what we are thinking and how things perceptually seem to us. The point I have been trying to make in this essay is that the master idea that guides his argument is a particular way of thinking, not about our knowledge of the contents of our own minds, but about our observational knowledge of ordinary empirical states of affairs. It is because he understands perceptual awareness of a red apple in front of one as he does that Sellars rejects a host of traditional ways of thinking about awareness of having a sense impression of a red apple or the thought that there is a red apple in front of one.

I have claimed Sellars understands the sort of perceptual awareness of external objects that is expressed in observation reports as the product of exercising two different sorts of capacities: the capacity reliably to respond differentially to stimuli (which we share both with merely sentient creatures such as parrots and with merely irritable devices such as thermostats and landmines) and the capacity to take up positions and make moves in a game of giving and asking for reasons. I have rehearsed the way I see some of the major arguments and conceptual moves in the essay as rooted in this two ply conception: the account of the use of 'looks' talk that underlies the incorrigibility of sincere contemporaneous first-person reports of how things

perceptually seem to one, including the treatment of scoped and generic 'looks' claims, Sellars's approach to the issue of concept acquisition, which caused so much trouble for traditional rationalists, and his rendering of the distinction between theoretical and observational concepts.

I would like to close with some observations and questions about the relations between the two kinds of ability whose cooperation Sellars sees as required for observation. The two sorts of capacities define dimensions of perceptual awareness that are in a certain sense orthogonal. We saw in the discussion of concept acquisition the broad outlines of a story about how one might move from possession of mere RDRDs to the capacity to apply observational concepts. And we saw in the discussion of theoretical and observational concepts how one might move from the purely inferential capacity to apply a concept, by the addition of suitable RDRDs, to mastery of a fully observational concept. That is, we saw in the case of particular observational concepts how one could have either of the two components without the other, and then move to having both.

But this shows only *local* independence of the two components: that one can have the RDRD of an observational concept without having the concept, and one can have a concept without having the RDRD needed to be able to apply it observationally. The corresponding global independence claim is not true. Purely theoretical concepts do not form an *autonomous* language game, a game one could play though one played no other. For one must be able to respond conceptually to the utterances of others in order to be talking at all. So one could not play the game of giving and asking for reasons at all unless one could apply at least *some* concepts noninferentially, in the making of observation reports. But this does not mean that

there could not be an *insulated* region of purely theoretical concepts, say those of pure mathematics—'insulated' in the sense that they had no inferential connection to anything inferentially connected to a concept that had an observational use. I don't say that any actual mathematics is like this, though it may be. Pure mathematics, I think, is in principle *applicable* to ordinary empirical objects, both those accessible through observation and those (now) accessible only inferentially. Applying an abstract mathematical structure to concrete objects is using the former to guide our inferences concerning the latter. But this relation ought not to be assimilated to that between theoretical objects and observable objects. It is not clearly incompatible with a kind of inferential insulation of the game of giving and asking for reasons concerning the mathematical structures. I think there are many interesting issues in the vicinity that are as yet not fully explored.[10]

It might seem that there could be no interesting question concerning the potential independence of RDRDs, corresponding to this question about the potential independence of the game of giving and asking for reasons. For it seems obvious that there can be reliable differential responsive dispositions without conceptual capacities. That is what mere sentients and artifacts have. But I think in fact there is a subtle question here, and I want to end by posing it. To begin with, what is obvious is at most that the RDRD's corresponding to *some* observational concepts can be exhibited by creatures who lack the corresponding concepts. And we might doubt even this. The story of John in the tie shop reminds us that our dispositions

[10] See for instance McDowell's discussion in "Brandom on Inference and Representation*," Philosophy and Phenomenological Research* Vol. LVII, No.1, March 1997, pp. 157ff, and my reply at pp. 189ff.

actually to call things red can be quite complex, and interact with our background beliefs—for instance about what are standard conditions for observing red things, and what conditions we are in—in complex ways. Though this claim goes beyond what Sellars's says, I think that learning about systematic sources of error can lead us to alter not just how we express our dispositions (substituting 'looks φ' for 'is φ'), but eventually even those dispositions themselves. I think, though I cannot say that I am sure (a condition that itself ought to give some sorts of cartesians pause), that familiarity with the Müller-Lyer illusion has brought me to a state in which one of the lines no longer even *looks* to me to be longer than the other. The more theoretically laden our concept of standard conditions for some sort of observation are (think of the mu-meson case, where those conditions involve the presence of a cloud chamber), the less likely it is that a creature who could deploy no concepts whatsoever could master the RDRDs of a sophisticated observer.

Besides creatures who lack concepts entirely (because they are not players in any game of giving and asking for reasons), we could ask about which RDRDs are in principle masterable by concept users who for some reason lack the specific concepts that for the genuine observer are keyed to the RDRDs in question. It might be, for all I know, that by suitable reinforcement I could be trained to sort potsherds into two piles, which I label with the nonsense terms 'ping' and 'pong', in such a way that I always and only put Toltec potsherds in the 'ping' pile, and Aztec ones in the 'pong' pile. What would make my noises *nonsense* is that they do not engage inferentially with my use of any other expressions. And we might suppose that I do not have the concepts Toltec and Aztec. If told to substitute the labels 'Toltec' and 'Aztec' for 'ping' and 'pong', I would then be a kind of idiot savant with respect to the noninferential applicability of

those concepts (which I would still not grasp). Perhaps there are no conceptual limits to such idiot savantry. But I find it hard to conceive of cases in which someone who lacks all the relevant concepts nonetheless can acquire the RDRDs necessary to serve as a measuring device (not, by hypothesis, a genuine reporter) of observable instances of the applicability of thick moral concepts such as courage, sensitivity, cruelty, justice, and so on. Of course, unless one endorses something like Sellars's account of what is required for something to be observable, it will seem that such properties are not suitable candidates for being observable by *anybody*, never mind by idiot savants. But for those of us who do accept his approach, this sort of question is one that must, I think, be taken seriously. That is the thought I want to leave you with at the end of this chapter.