

CHAPTER SEVEN INTENSIONALITY

By stressing that all complete signs signify complete world affairs in the previous chapters, I may seem to have implied that complete signs are always translatable by sentences. To see why this is wrong we need to understand how signs are used to represent other signs. For the difficulty lies in the fact that the only direct way we have to speak of what non-sentential signs represent is by misleadingly comparing them with sentences.

In Chapter Four I discussed the way defining descriptions work --by intentionally representing natural signs. Intentional signs may also be used to represent other intentional signs. Sometimes this produces the phenomenon philosophers call "intensionality" (with an "s"), as I will explain. There are also some cases in which intentional signs are used to represent natural signs that are not as straightforward as the cases of defining descriptions. These can give rise to a different (perhaps previously unrecognized) form of intensionality.

Wilfrid Sellars claimed that the form of expression "'X' means Y" --as in "'Hund' means dog" or "'rouge' means red" or "'Chicago est grande' means Chicago is large"-- does not assert a relation between an expression and some other entity, say, a property or a world affair. To understand the meaning of the "means rubric" is to understand what its characteristic purpose or function is. Its function is to produce in the hearer a disposition to use the expression "X" in the same way that the hearer already knows to use the expression "Y" in his/her home language. Thus, so long as the two expressions matched up in the "'X' means Y" formula play, as Sellars put it, "the same role" in their respective languages, the form "'X' means Y" is used correctly. The beauty of Sellars' account is that it works just as nicely for expressions that obviously are non-denoting, as in "'et' means and," "'arret!' means stop!," "'Hélas!' means Alas!" --or "'signifier' means mean"-- as it does for "'Hund' means dog." On the other hand, taken by itself, this account does nothing to clarifying what in the world a "linguistic role" is. It portrays the expression "'X' means Y" as a translation rubric but, taken alone, it tells us nothing about what it is for one expression to be a good translation of another.¹

Similarly, in "On Saying That," Davidson (1968-9) claims that a sentence such as "Galileo said that the earth moves" is true just in case uttering the words inside the "that..." clause of this sentence makes the speaker and Galileo into "same sayers." But again, we are left wondering just what it is for two speakers to say the same thing. It seems reasonable to see Davidson's account as strictly comparable to Sellars's, however. To be "same sayers" is to use expressions that have the same "linguistic role" --whatever that is.

Davidson's analysis is easily applied also to the expression "...says to... ." In asserting "Mother says to wear your leggings" the speaker claims to be a samesayer with Mother. Mother has said words having the same linguistic role that "wear your leggings" would have if actually used, rather than merely displayed, in the present context to the present hearer. Davidson's analysis is also applicable to forms such as "...believes that...," "...intends to," "...wishes that..." and so forth, on the assumption that beliefs, wishes, intentions, and so forth are mental representations and that a mental representation can, in some sense, "play the same role" as a linguistic representation. That, I believe, was how Sellars saw the matter. He could do so because "playing the same role" was, for him, very much a matter of more or less (1963, chapter 6). Both men suggest that we represent a representation by holding up another representation that is

similar to it in relevant ways. I will argue that which ways are relevant ways is generally determined pragmatically rather than being grammaticalized and that this results in intentionality.

I have drawn attention to language forms --phonological structures, words, syntactic forms, aspects of prosody and so forth-- as reproduced entities (Chapter Two). They are memes, with natural purposes that may differ from or "cross over" the immediate uses to which individual speakers put them. Given the description of intentional representation offered in Chapter Six, public language forms (types) are intentional representations just when fulfillment of their functions or purposes by normal mechanisms, which entails the collaboration of trained cooperative hearers, requires that they coincide with affairs in the world according to established semantic mappings. Their continued reproduction has depended their having served cooperative purposes of speakers and hearers often enough, and this has depended in turn on correspondence between them and world affairs onto which they have mapped by rules to which speakers and hearers are both adjusted. Broadly, then, my suggestion is that the best sort of translation of a language form will match both its purpose, that is, its linguistic or memetic function, and also its semantic mapping function (note the two different senses of "function"). Generalizing this, the best sort of translation of any intentional sign will match both its purpose and its semantic mapping function.

If this is right, the possibility of saying (showing) precisely what an expression in another language "means" by the method of samesaying will depend on the availability in one's home language of an expression having both a matching purpose and, if the expression is an intentional representation, also a matching semantic mapping function.

Consider, for example, the intentional signal the rabbit produces when its predator-detectors fire. The rabbit thumps its hind feet smartly on the ground. The natural purpose of this is to trigger a reflex that causes its relatives (rabbits, not just any bystander) to freeze or take cover. There is no literal translation of that particular pushmi-pullyu danger-thump into English or French. "Danger!", for example, tells of danger to humans, and "Rabbit danger!" does not have the function of sending any rabbits to cover. Rather than directly saying "the rabbit thump means ...(so and so)," to achieve accuracy we must dispense with the "'X' means Y" formula and set to describing its purpose --as I did just above. Similarly, in trying to explain what the formula "'X' means Y" itself means, Sellars didn't use samesaying. He described what its function is, what reaction it is used to cause in a hearer.

Turning now from linguistic function to semantic mapping, the linguistic function of the "'X' means Y" formula will be performed in the normal cooperative way only if "X" has the same linguistic role in its language that "Y" does in its. Saying "'X' has the same linguistic role as 'Y'" is thus a way of stating truth conditions for that formula. But it does not follow, for example, that the expression "'Hund' means dog" means the same as the expression "The word 'Hund' plays the same role as does the word 'dog'." For these expressions do not have the same linguistic function. Very small children do not have concepts of words, as fully understanding the latter sentence would require them to have. For example, they will claim that since there are no ghosts, "ghost" is not a word (Susan Carey, private correspondence). But they acquire the ability to react appropriately to the verbal "X means Y" formula very early on. Similarly, the truth condition of a sentence asserting identity, "A is B," requires the word "A" and the word "B" to have the same

referent or extension, but understanding the sentence does not require thinking about words (Millikan 2000, Chapters 10-12). The job of an indicative sentence is not always to cause a belief with the same truth condition that the sentence has (Millikan 1984 Chapter 12, 2001b). The functions of intentional signs can come apart from their satisfaction conditions.

Besides the difficulties that concern matching at once both linguistic function and semantic mapping function, there may be difficulties in matching semantic mapping functions just taken alone. For there can be crucial differences in articulation of representations that have, none the less, identical truth or satisfaction conditions. Semantic mapping functions are not the same things as truth or satisfaction conditions, and can easily come apart from them. This is because a sign is, essentially, a member of a system of signs, and the same piece or aspect of the world --the same truth-maker-- can be represented in sign systems that are not isomorphic to one another. I have said there is no translation of the rabbit's danger thump into English or French because these languages contain no forms with the same primitive function. Another reason is that the semantic mapping function that aligns rabbit thumps with the affairs that satisfy them does not articulate these affairs in a way that parallels the semantic mapping function for any English or French sentence. The rabbit thump sign has exactly two variables, time representing time and place representing place. Rabbit thumps are articulated the way stop lights are. Move the time and place of the red light and that moves the time and place to stop. Contrast a rabbit thump or a currently lighted red light with the sentence, "Stop here now!" The sentence is articulated so as to contrast with "Stop over there now" and "Stop over there in an hour," and with "Sit here now" and "sit over there tomorrow," and so forth. It is also subject to a negation transformation: "Don't Stop here now!." The rabbit thump is not a member of any such system of signs. No transformations of it tell of times other than its own time, or of other places, or of things other than rabbit danger, or tell when or where there is no rabbit danger. So to say "Freeze or take cover here and now!" is not to samesay accurately with the rabbit. Would "Freeze or take cover!!" come closer to samesaying with the rabbit (ignoring that its function is not to affect rabbits but humans)? But "Freeze or take cover!" contrasts with just "Freeze!" or just "Take cover!," and also with "Hop or jump!". More interesting, there is no negation transformation of the rabbit thump parallel to "Don't Freeze or take cover!". Nor is the rabbit thump transportable into other contexts, as the word "Freeze" is in "Freeze when I blow the whistle!" Still, the satisfaction conditions of the rabbit thump are expressible in English. The indicative satisfaction conditions are that there is danger to rabbits at the time and near the place of the thump. Those are the thump's truth conditions. It's imperative satisfaction conditions are that the rabbits now freeze or take cover. To express satisfaction conditions you do not have to samesay. But satisfaction conditions do not reveal semantic mapping functions. To tell what the truth conditions are is not to reveal the significant articulation of the representation.

Compare the dance of the honey bee. It represents the current location of nectar relative to the bees' hive and the direction of the sun. But there are no transformations of it that would tell about nectar location relative to objects other than the hive and the sun, or about the location of anything other than nectar. Its references to the nectar, the hive and the sun are all implicit. Only the reference to the angle between the nectar and the line from the hive to the sun is explicit. No

English sentence with the same truth conditions approaches this degree of inarticulateness. I cannot tell you where the nectar is relative to the bees' hive and the sun without explicitly mentioning at least nectar, explicitly mentioning or describing the relevant hive, and explicitly mentioning the sun. An English sentence with the same truth conditions is subject to significant transformations that will tell instead of the relation of nectar to hive and moon, or of nectar to the Eiffel Tower and the moon, or the relation of peanut butter to hive and sun, and so forth. Nor is the language of the bees "systematic," in the way Fodor and Lepore use that term: "If... a language can express the proposition that aRb, then it can express the proposition that bRa" (1992, p. 146). The bees have no way of saying that the sun lies at a certain angle between the hive and the nectar. It is true that bees sometimes represent with their dances the location of water, if water is much needed or, when they are swarming, of suitable places to build a new hive. But there is nothing in the bee dance itself to indicate this shift to a different semantic mapping function. The shift is recognized in the same sort of way that the domain of a local sign (or of a defining description) is tracked. The bee has to independently "know," as it were, which local domain this sign is in, for the dance itself doesn't say. (The bee has to understand pragmatics as well as semantics.) A sentence that represents the location of nectar relative to the sun and the bees' hive will also be subject to a negation transformation. It will contrast with a sentence representing that there is no nectar there. But bees have no way of saying where there isn't any nectar, so don't bother looking. Other bee dances tell of nectar at other places, but to say there is nectar one place does not contradict that there is nectar another.

Having distinguished between semantic mapping functions and satisfaction conditions, we can apply this to the philosopher's notion "proposition," which hovers between these two. Volumes could be written about this ambiguity and the trouble it has caused, but let me just cite one example. Fodor and Lepore say the sentences of a language are, in general, "isomorphic" to the "propositions" they express: If a sentence expresses the proposition that John Loves Mary then there will be elements corresponding to John, to loves and to Mary (1992, p. 147). In a footnote they then remark that this is debatable for the sentence "It's raining," which they take to correspond to a proposition about place and time as well as rain. And they say that this seems not to be true of certain idioms as well. Having thus waded in and barked their shins, they complain that the waters in this area are muddy. What they are stumbling into here is the ambiguity in the notion of a "proposition." Semantic mapping functions are different from truth or satisfaction conditions, but the notion "proposition" hovers between, sometimes coming to rest on one side and sometimes on the other. The notion of "the proposition expressed" presupposes that semantic mapping functions, which are determined by "compositionality" in the broad sense, that is, by "architectural structuring" (Chapter Four), are the same as truth or satisfaction conditions --But they are not.

Also consider Evans' "generality constraint" in this connection. Evans held that in order to think of a thing it is necessary to know what one is thinking of, that this requires that one have a "concept" or "Idea" of that thing, and that a concept or Idea is a general ability that "makes it possible for a subject to think of an object in a series of indefinitely many thoughts, in each of which he will be thinking of it in the same way" (Evans 1982 p.104). That is, it is not possible to think of a thing unless one can represent it as embedded in many alternative kinds of states of

affairs. Something like this may be right as a requirement for having a concept of something (Millikan 2000, chapters 13 and 14), but it would be a mistake to claim that this was a necessary condition on all inner or on all intentional representation. The bees, for example, surely represent that there is nectar at a certain angle off the line between sun and hive without expressing detachable concepts of the nectar, the sun and the hive. Moreover, inner pushmi-pullyu representations such as thirst, hunger, and pain represent occurrences of inner states at times and direct appropriate action without further articulation. Thirst, for example, is not an articulate desire for water, but its truth condition is that the body needs water. It directs drinking, but it is not an articulate directive like "You drink water now!"

Thus it is that the public-language representations of humans, and surely also human beliefs, desires and intentions, may differ quite radically from more primitive inner and outer representations having the same satisfaction conditions, representations that are either used by humans below the level of explicit belief, desire and intention, or used by other species. Descriptive sentences in all human languages have, at a minimum, a subject and a predicate and are sensitive to a negation transformation. These properties set them far apart from a host of simpler representations, undoubtedly including many kinds of inner representations that help to govern human behavior on levels lower than that of rational thought. An important part of the story of the evolution of cognition, with which I will be concerned in Part Four, will concern the emergence of various new forms of articulation, as well as new functions, for inner intentional representations.

It is on this dimension that sentences differ also from other sophisticated intentional representations such as maps, charts, graphs and diagrams. The correctness of an ordinary map or diagram may entail the truth of various sentences. But the affairs mapped in common are projected by semantic mapping functions that articulate these affairs against quite different contrasting possibilities. The space of significant transformations surrounding these different kinds of representations is entirely different. Each resides in what early Wittgenstein might have called a different "logical space". And, of course, this dimension is also often relevant when comparing different linguistic expressions to one another. In Philosophical Investigations, Wittgenstein remarks on the difference between saying or meaning that the broom is in the corner and saying or meaning that the brush and the stick are in the corner. Or consider the difference between believing you are drinking water and believing you are drinking H₂O.

Having come this far we can more easily understand the phenomenon of intensionality. Sellars and Davidson each pointed out in his own terms that we represent signs by displaying other signs that are like them. That is, we "portray" or offer portraits of signs we wish to talk about by holding up similar signs. Notice that this does not make the portraying sign into an indexical. Its kind stands for another of the same kind just as the place of quail tracks stands for quail in the same place and the size of the tracks stands for quail of the same size. The kind is a merely a "reflexive" element of the sign (Chapter Four). But almost always, the portrait is like the original only in certain respects -- respects that happen to be relevant to the communicative purposes of the moment.

First, the possibility of using samesaying to express precisely the properties of a sign that is not in one's home language depends, as I have said, on the availability in one's home language of an expression having both a matching semantic mapping function and, if the sign to be

represented is an intentional sign, also a matching purpose. But as Sellars pointed out, "playing the same role" is very much a matter of degree, and matches that are far less than perfect are often perfectly serviceable. Indeed, often only certain properties of the sign to be portrayed concern us, so that the portraying sign needs to be like the portrayed sign only in very limited respects. Second, there are times when we want to convey information about more than the role of a representation. Sometimes the very words the speaker used make a difference. Then samesaying may require using the very same words as the speaker, or words that are like them in relevant physical or etymological respects. But unfortunately, when one sign is held up to portray another, which respects of likeness are the relevant ones on the given occasion is not generally, as linguists put it, "grammaticalized." One usually relies on pragmatics to be sure that the hearer understands what aspects of the sign held up are the ones being attributed to the represented sign.

It is customary to describe intensional contexts as contexts in which coreferential terms can not be substituted for one another without change of truth value. For example, although Bernard J Ortcutt may be the same man as the man that Ralph has seen in the brown hat, since Ralph may not know this is so, "Ralph believes that Bernard J Ortcutt is a spy" may have a different truth value than "Ralph believes that the man in the brown hat is a spy." This is supposed to show that "Ralph believes that..." is an intensional context. But in fact it is often possible to substitute coreferential terms inside "...believes that..." contexts with no risk of changing truth values. This is because the purpose of holding up a portrait sentence within a "believes that..." context is often merely to portray aspects of the reference of the believer's thought. For example, if I say to you "Ralph thought that your venerable dean was a spy," the fact that Ralph has no idea that either Bernard J. Ortcutt or the man in the brown hat is a dean, let alone your dean, has no effect on the truth value of my sentence. The phenomenon here is not that one cannot substitute coreferential terms without change of truth value, but that the grammar alone does not prove that one can. Whether one can or not is a pragmatic matter. I suggest that an intensional context, described in a more general and illuminating way, is merely a context in which one sign is held up to portray another but where grammar alone does not tell what kind of likeness is intended. Consequently, grammar alone does not tell what other signs might be substituted for the sign held up without altering the import. The phenomenon concerns not merely substitution of coreferential terms but various other substitutions as well, such as substitutions of single words or of phonemes or of words with different etymologies.

Consider "John said the earth moves." The most typical reading of this sentence would be to take "the earth moves" as a portrait of the semantic role of the sentence that John uttered. "John said the earth moves" would then be taken to convey about John the same thing that "Galileo said the earth moves" truly conveys about Galileo, even though Galileo did not speak English. But "John said the earth moves," uttered by anyone, or written by a someone who does not care about philosopher's conventions with quotation marks, could also be used to portray John's very words, the vehicle as well as the semantic role of his representation. Indeed, it is even possible to use that sentence to portray the vehicle of John's representation only, forget it's role. Suppose that I am aware that my hearer does not understand much English. Perhaps John and I are just helping her with English phonemes, and I am just repeating John's sentence for her with clearer pronunciation. Nothing in the form of the sentence "John said the earth moves" shows

which of these interpretations is intended, a portrait only of phonemes, or a portrait of certain definite meaningful words, or a portrait only of semantic content, or some combination of these.

If I say instead, "John said that the earth moves" perhaps the grammar indicates that the embedded sign I hold up portrays only aspects of the meaning, and not of the vehicle of the sign talked about. I may say it if John spoke German or Italian rather than English. Yet this is not always the case either. Consider "John kept insisting that there were many more Greeks than Hellenes." Given that the Greeks were the Hellenes, here it must be that the sign vehicle John used is being portrayed too. The situation that would normally be portrayed by the "that ..." clause in "John insisted that there were many more Greeks than Greeks" is quite different. Result? The phenomenon called "intensionality." There are contexts in which single words having the same referent cannot be exchanged without "changing truth value," because exchanging them changes what kind of sign would be taken to be portrayed, a sign typed by vehicle or a sign typed by meaning. That the context is one of this sort is not a matter of the grammar alone, but is a pragmatic matter.²

Now compare "John said that your mother is brave" with "John said that the governor is brave," as said by Tom to little Willie. Assume that Willie's mother is the governor. Recall that defining descriptions such as "your mother" and "the governor" can be used by speakers to serve any of three functions (Chapter Four). It may matter or it may not matter whether the hearer knows of whom the properties mentioned in the description are natural signs, and it may matter or it may not matter whether the hearer understands and keeps in mind the properties mentioned. Now unless John, the original speaker, had been speaking to Willie himself, presumably he would not have used the words (the vehicle) "your mother." It is pretty clear pragmatically, then, that John's exact words are not being portrayed. So what is being portrayed? Did John refer to the governor as Willie's mother, or not? Likewise with "the governor." Did John refer to Willie's mother as being the governor of something or not? Whether the sign used to portray the intentional sign John used follows the same "route" (Chapter Four) to the further affair represented by John's sign or whether it doesn't is not shown in the grammar of these sentences. But it may well be clear from context what aspect of the sign John used was central to John's purpose in using it, or what aspect is central to what the current speaker wants to convey. For example, although John didn't refer to the governor as being Willie's mother, it may be important to Tom, the current speaker, that he convey to little Willie exactly that it is Willie's mother who is admired by John. That may make little Willie proud.

More generally, unless the pragmatic context clearly indicates otherwise, names and defining descriptions that appear within "said that..." contexts are not usually taken to portray either the vehicles or descriptions used by the original speaker. Rather, they portray more distal affairs that are signified by the vehicles or by the properties represented on the route to these distal affairs. The names or descriptions that are used are chosen to convey these distal affairs in whatever way the current listener will most easily understand within the conversational setting. But again, it is clear that there do exist some contexts in which coreferential terms cannot be exchanged without changing truth value under at least some pragmatic circumstances. They cannot be exchanged without changing what kind of sign a sentence held up as a portrait would normally be taken to portray. These are intensional contexts.

Since the memetic purposes of public language forms are never exactly the same as the natural purposes of inner representations such as perceptions and thoughts --obviously they do not do exactly the same jobs-- no exact translation of what any inner representation "means" can be given in a public language. It is impossible strictly to samesay with someone's belief. That may be one reason why there is something unnatural about saying that thoughts and perceptions "have meanings," even on the assumption that they are inner representations. You can't say what they mean in a straightforward way. Certainly you couldn't portray exactly what they meant taking their exact functions into account. But there is another and deeper problem about representing thoughts. There has been considerable controversy recently about whether ordinary language embodies a sort of "folk theory" about thoughts, a theory that implies that thoughts are representations in people's heads that can be fairly accurately portrayed using ordinary sentences. Advocates of this view usually claim that the folk also believe that our behaviors result from causal interactions, corresponding to inferences, among these inner sentence-like representations.³ Happily, I don't need to be concerned here with how "the folk" explain behavior. But it does seem clear that the way we talk about our intentional mental states does, rightly or wrongly, assume that these states are enough like sentences to be portrayed by holding up sentences. In the second half of (Millikan 2000) I argued at some length that sentences may well be a very misleading model for thoughts. But if we want to understand only how representations of thoughts are employed in ordinary language, we can withhold judgment on what thoughts are actually like. We can examine, merely, how they are portrayed by ordinary language, given our use of expressions such as "believes that...," "fears that...," "intends to...," "hopes to..." and so forth.

They are portrayed as being much like inner sayings. Typically, I believe, they are thought of as analogous to sentences the thinker would use to express these thoughts candidly.⁴ And the way "...says that..." and "...says to...," are used corresponds very closely to the way "believes that...," and "desires that" and "... intends to...," and so forth, are used. Although the sentences held up as portraits in these contexts are intended to bear a resemblance to sentences the speaker might candidly have used, just how close a resemblance depends on the pragmatic context. In particular, defining descriptions (Chapter Four above) captured inside these contexts may be intentionally used to portray descriptions the thinker in question would himself have been disposed to employ, or they may be used merely referentially.

The default assumption is probably that what is portrayed is the reference. The descriptions are chosen mainly for the purpose of getting the hearer to understand the references, given the hearer's background and the current pragmatic context. Descriptions the original thinker might have used to express his thought, should these even be known to the speaker, probably play a role only occasionally, when it happens that something more turns on them. So although contexts in which intentional attitudes are portrayed are usually contexts in which coreferential terms can be exchanged without changing truth value, this is not always the case. Under some pragmatic circumstances, changing the descriptions will change the understood truth value. The ancients believed that the morning star was seen in the morning but they did not believe that the evening star was seen in the morning. Contexts in which intentional attitudes are portrayed are intensional contexts.

There are even circumstances in which the sign held up inside the "that..." clause

following an intentional verb portrays its own vehicle as one the thinker would use to express his thought. Using the same example again, it is clear what "John firmly believes that the Greeks were more numerous than the Hellenes" has to mean, at least in a context where speaker and hearer both know that the Greeks were the Hellenes. It has to mean that John would have expressed his thought using the very words "the Greeks" and "the Hellenes." Compare also "John was quite sure that a new spigot would cost more than a new faucet!" Given that spigots are the same thing as faucets, it is pragmatically clear that John was a speaker of English, or that he at least had these two words of English. (Even if one believes in Fregean senses, it is completely implausible to suppose that the public words "spigot" and "faucet" express different Fregean senses! Equally, I suggest, for "John was quite sure that Cicero was born before Tully.")

Modal contexts are intensional contexts. I would like to defend the view that modal contexts too are best analyzed as containing representations of representations. Talk of possible worlds is just disguised talk of representations. Though I think this view is correct, the argument for it will have to wait for another occasion.

Descriptions of causes, of natural explanations, and of natural purposes are sometimes cited as creating intensional contexts. But this, I believe, is an error. Consider, for example, the difference between "What caused the fire was that Herbert's youngest child was playing with matches" and "What caused the fire was that Billy was playing with matches." That the child playing with matches belonged to Herbert surely was not relevant to causing the fire. But this merely exemplifies Donnellan's distinction again (Chapter Four). A hearer has to gather from context whether the speaker's purpose is, or is in part, to convey information about the properties mentioned in the description or whether these properties are to be understood as relevant only as a sign of something else (the referent) about which information is offered. Talk about causes is not talk about signs. Similarly, descriptions of natural purposes are subject to Donnellan's distinction.

In this chapter I have discussed phenomena that give rise to linguistic contexts that are "intensional" (with an "s") in the modern classical sense. Notice that there has been no reference to any relatives of Quine's "creatures of darkness," to "intensions" as understood in the tradition of Carnap, nor to any relatives of Fregean "senses." I have been describing only differences in purposes and differences in semantic mapping functions that map either natural or intentional signs onto the extensional affairs that they signify. The "intensional" has been explained in completely extensional terms.⁵

FOOTNOTES

1. See, for example, Sellars 1963, especially Chapter 6.

2. I do not take predicates having the same extension to be the same kind of phenomenon as descriptions having the same referent. I assume a realism about properties. My remarks about coreferential descriptions should not be taken to generalize to merely coextensional predicates.

3. It is often assumed, further, that the folk take these causal interactions partly to determine the very meanings of these inner representations. But this last description of how the folk think about thought is implausible, given the history of philosophy of mind, for no one made this suggestion prior to the 20th Century.

4. It does not follow, I hasten to add, that ordinary people think they are "individuating" thoughts by reference to a one-to-one correspondence to possible

sentences. The thinker might have many different ways at her disposal of expressing the same thought, for example, many ways of expressing thoughts of the same subject by using different defining descriptions. There is no reason to attribute to the layman a philosopher's conflation of a way of recognizing something with a way of thinking of it. On this confusion, see (Millikan 2000) Chapters Eight and following.

5. Once again I must mention Wilfrid Sellars. He was as concerned as Quine to escape this kind of darkness, and my way of escaping has made use of paths that he originally forged.