

In The Externalist Challenge. New Studies on Cognition and Intentionality, ed. R. Schantz (Berlin & New York: de Gruyter).

EXISTENCE PROOF FOR A VIABLE EXTERNALISM

Externalism, as I am understanding the term, is a thesis about the nature of thoughts, as distinguished from language. For example, Kripke's suggestion that the referent of a public language proper name is determined by its history is not, just as such, an externalist thesis. On the other hand, Putnam, in "The Meaning of 'Meaning'," did seem pretty clearly to be talking about the nature of our thoughts of such things as water, beech trees and aluminum, perhaps as well as about language. And his essay is often taken as the original and also the paradigm defence of externalism.

Putnam argued that a person's internal psychological state does not determine the referent or extension of that person's thought. Hence, if what a person means or intends were determined solely by that person's psychological state, what a person means or intends would not determine the referent or extension of his thought. Denying the consequent, Putnam concluded that what a person means or intends with a thought is determined by more than that person's internal psychological state. As he put it, "Meaning just ain't in the head!"

If we explain the externalist idea in this crude way, however, it becomes hard to see how anyone could deny it. If the question were, merely, how are the referents or extensions of thoughts determined, it seems patently obvious that nothing inside someone's head could, by itself, determine that anything in particular existed outside the head. Referents and extensions are existent things and existent sets. What happens to exist or not exist, when and where, outside one's head is surely a contingent matter. How could what is inside a head determine that anything at all had to exist outside that head? But if not, how could it determine, all by itself, that its thoughts had referents and what these referents were? Something has been stated wrong. Externalism should not be so obviously true.

One's first thought here may be that Putnam has denied the wrong premise. What seems obvious is that meanings, taken alone, without adding the world, do not determine extensions. But how has it been shown that what is in one's the head does not determine what one means? Let us take a second look. What exactly did Putnam's arguments show?

What he seems actually to have argued is that what is in the head, when combined with what is in the world, does not determine the reference of a thought in the way that was classically supposed. He argued this, specifically, for the case of thoughts of natural kinds. Classically, it was supposed that what was in the head picked out a set of general properties, and that the world determined what, if anything, had these properties, and that if anything in the world did have these properties, it was part of the extension of the natural kind thought. Putnam argued that natural kind thoughts were, instead, "indexical." By this he apparently meant two things. First was that the extension was determined by some sort of concrete existential relation between the thought and (parts of) its extension, such as a causal, temporal and/or spatial relation. (Water is the so-looking and so-feeling stuff around here.) Second was that the relation between the thought and its extension was not determined by being represented in the mind. This relation was not determined by being thought of.

Now many have felt that there is nothing in Putnam's arguments to support the second of these theses. See, for an especially clear statement of this, (Fumerton 1989).

Be that as it may, surely there is nothing in Putnam's arguments, at least, to show that there is nothing at all within the head that determines which relation the extension-determining relation shall be. Putnam's slogan "Meaning just ain't in the head," if we take "meaning" to be whatever, taken along with the world outside the head, determines extension, does not seem to be supported by his arguments. Hence a very common reaction to Putnam's arguments has been, exactly, to claim that what is in the head determines what kind of relation must exist between a certain thought and its referent or extension for that to be its referent or extension, though not necessarily by representing that relation. Perhaps that relation is the relation of exemplifying certain properties or relations whose identity is determined by what is in the head, or perhaps it is the relation of being the cause of the thought, or perhaps it is the relation of covarying with the thought in a counterfactual supporting way, and so forth. Indeed, perhaps this reference determining relation can vary from one kind of thought to another.¹ Taking the meaning (as distinguished from the reference or extension) of the thought to correspond to this relation or to whatever determines it, Putnam's observations about natural kinds may be true, yet fail to show that meaning is not in the head.

Now I have no wish to argue over how the terms "internalism" and "externalism" should be used. Nor do I think that generalized arguments for or against abstract philosophical positions --internalism/externalism, realism/antirealism, individualism/anti-individualism and so forth-- are ever of much, if any, value. What may be of value, however, is to lay out a well-articulated position on how people's heads actually do interact with the world they are in to create the phenomena of meaning and reference in thought, and to show how this concrete position explains various phenomena that we are interested in, at the same time avoiding certain problems that we should wish to avoid, and that have concerned those interested in the internalism/externalism debates.

I have attempted to articulate such a position over the years in various books and papers. The position happens to be, in what seems a very strong sense, externalist. For it implies not only that basic reference or extension is always determined by a concrete existential relation between the thought and its referent or extension, but that what determines this relation to be the relation that determines reference or extension is not merely a matter of what is currently in the head. This claim might reasonably be expressed, using Putnam's phrase, by saying that meaning is not in the head. But it is a more extreme position than the one Putnam actually gives us arguments for in "The Meaning of 'Meaning'."

I cannot, of course, re-present the whole of a developed theory of thought reference and thought meaning in a few pages. But what I can do is to outline the theory, supply references for where more complete exposition and defense can be found, and then concentrate on showing how this particular theory addresses basic concerns about externalism that have been voiced by internalists.

The theory takes thoughts to be inner representations. It is peculiar in what it takes inner representations to be. One way to understand it is as a modification of classical functionalism. It modifies classical functionalism in two ways. First, it defines inner representations by the way they function, not just in the head, but as parts

of much larger systems that include portions of the environment.

Second, the functions by reference to which inner representations are defined are not mere dispositions of the representations within their inner and outer environments, but what I call "proper functions." These are dispositions they were selected by natural selection for having, or dispositions that a Normal development of the biological system has produced by means of the organism's interacting with its environment in a Normal way.

I capitalize "Normal" to mark off a special sense of that word. What is biologically Normal, as I use that term, is not what is common or average, but the way examples of that lineage of biological systems have generally functioned in the past on those occasions that accounted for their selection --better, that accounted for these lineage members not having been selected against in situations where members of the lineage not functioning this way were or would have been selected against.² It is this reference to a certain kind of history of selection and/or development that adds the radically externalist twist to this theory of mental representation. What a thing was designed to do is not always evident just from its inner structure, even from its inner structure plus the structure of its current environment. Accordingly, whether an inner happening or structure is a representation is not merely a matter of its inner structure.

Inner representations are defined by reference to the way representations Normally function in a wider biological system that includes the organism's Normal environment. But they cannot be defined merely as items that are in fact functioning in a Normal way. Rather, they are defined as items produced by systems --genetic systems, perceptual or cognitive learning systems-- that, if functioning Normally, would produce representations that were capable of functioning Normally, given Normal operation of the rest of the inner cognitive systems, and given the actual outer environment within which these systems are operating. But these producing systems may not have been functioning Normally. In that case, we may sometimes say that what they produced were indeed inner representations, but not Normal ones. Primary ways for representations to fail to be Normal are being false, or empty, or equivocal. I will come to the full story on that a bit later.

At the moment, an important point I want to make is that when representation producers do not function Normally, this is not, in general, because there is anything wrong with these producers. Usually it is because the environment they are operating in is not Normal, relative to their particular historically Normal method of functioning properly. This point needs to be leaned on, since failure to understand it has produced well-known but mistaken criticisms of the position I wish to advocate (Dretske 1986, 1988; Neander 1994). Let me illustrate the point with an analogy.³

There are such things as ice-cream-making machines. Their function is to make ice cream. But this is a function they cannot perform, no matter how well they are built, unless they are in an environment that supplies certain materials and conditions. The right ingredients have to be placed by someone or something into the right input containers. And, in the usual case, the machine has to be right side up and plugged in to a source supplying the right kind of electrical power. And the machine has to be turned on. So much is trivial. But the obvious conclusion has often been overlooked. If a device is failing to perform its function, this does not necessarily imply that the device

itself is malfunctioning. It does not imply that the device is sick or broken. Mostly, when devices are not performing their functions, it is because the environments they are in are wrong, or the necessary inputs are not being put in or, for the most trivial example, because they are not turned on.

Similarly, eyes that are failing to see may fail because they are in the dark or because they are closed. Eyes that are seeing double (not something, presumably, that they were selected for doing) may do so because the optometrist has interposed an instrument between them and what they are seeing that constitutes an abnormal environment relative to their historically Normal way of functioning properly. Similarly, eyes that see mirages are trying to see under atmospheric conditions not historically Normal for distance vision. And so it is for systems whose proper functions are to produce non-empty, nonequivocal, true representations. They only work properly and Normally when placed in the right sorts of environments and given the right sorts of inputs.

I have said that the theory of thought that I advocate defines inner representations by the way they function, not just in the head, but as parts of much larger systems that include portions of the environment. Like systems that use inner representations, ice cream machines require to be in the right sort of environment and to have the right sort of input in order to work properly. But they are not defined by reference to the way they function, but only by reference to the output they are designed to produce. In what sort of way, then, is an inner representation designed to function?

I use the term "representation" in much the same way as the zoologist, C. R. Gallistel:⁴

I use the term representation in its mathematical sense. The brain is said to represent an aspect of the environment when there is a functioning isomorphism between an aspect of the environment and a brain process that adapts the animal's behavior to it. (Gallistel 1990, p. 3)

Unpacking this a bit, the idea is that it is, in part, because the animal responds to its environment with behavior-governing brain states that have aspects that run isomorphic to certain aspects of that environment that the animal's behavior manages to vary as a function of the environment so as to be adapted to that environment. Behavior needs to vary as a function of the environment. Inner representations supply a mechanism by which that is accomplished.

It is important that this description of inner representations requires them to have behavior-governing functions. For example, inner representations are not representations merely because they covary with environmental causes. The representation must, Normally, be used by the biological system as a representation. The isomorphism between representations and environmental aspects has to be one that helps to explain how the animal's behaviors become adapted to its environment so as to serve their biological functions or purposes.

Unpacking this idea even more, a Normally operating inner representation is part of a total system that includes (1) some aspects of an organism's environment, (2) a representation producing mechanism, and (3) a representation using mechanism. The producing mechanism produces representations that run isomorphic to the

environmental aspects as defined by some rule of correspondence. The representation-using part of the system relies on this particular rule of correspondence in producing behaviors that will be effective given that the environment actually displays the aspects represented.⁵

This use of the term "inner representation" is very broad indeed. Myriads of things inside the body other than thoughts are representations in this sense. But the internalism/externalism issue primarily concerns thoughts, beliefs and so forth. How then, exactly, do thoughts and beliefs fit in?

Although I have argued elsewhere that sentences can be a very misleading model for beliefs,⁶ in the current context that model will serve us well enough. We can think of beliefs as like mental sentences, containing words that are rearrangeable into other mental sentences that mean different things. We can think of beliefs as being compositional, and think of the extensions of the true ones --their truth makers-- as determined by some sort of Tarskian mapping. Certain definite kinds of rearrangements of sentence parts correspond to certain definite kinds of possible transformations on the aspects of the world to which the sentences must correspond to if they are to contribute to Normal functioning on the part of their users. Certain substitutions of sentence parts correspond to substitutions in aspects of world affairs mapped and certain additions correspond to additions, and so forth. That is, there is a functional isomorphism between the set of possible mental sentences a certain person could affirm and the set of world affairs to which these sentences would correspond were that person's cognition to proceed entirely Normally.

But cognition proceeds perfectly Normally only in the case of a person all of whose beliefs are true. (This last underlines the importance of distinguishing what is Normal from what is common or average.) In the case of beliefs, presumably Normal function involves, at least, interaction with other beliefs to create new beliefs, interaction with desires to produce more desires, and so forth, a final result being decision-making and productive action.

The semantic rules that define the Normal correspondences between a person's belief states and aspects of their environment are determined mainly by the history of that person's individual cognitive systems. Let me explain.

Mechanisms whose proper function is to produce inner representations that will correspond to the world in accordance with certain definite semantic rules have to have a systematic method or methods of achieving this goal. This does not mean that the methods always work, of course, that no mistakes are ever made. It means only that there must be conditions under which the method or methods do work, these conditions having been present in the past on the occasions when mechanisms worked well enough to be retained or selected for (rather than being selected against). Some mechanisms have been selected for by natural selection during evolutionary history. Others are selected for or tuned for their jobs through processes of learning. Mechanisms whose proper functions are to turn out beliefs that correspond, in accordance with definite semantic rules, to the world, have probably invariably been selected for or tuned by learning. What kind of learning is this, and how does it take place in a Normally developing cognitive system?

To supply details here is surely, in the end, a job for developmental

psychologists and neurologists. But some part of the story is discernable even at the distance of the philosopher. In order to learn to make beliefs that correspond in systematic ways to the world, one must learn how to identify, through the senses, the various objects, kinds, properties and so forth to which the elements of these beliefs are to correspond. That is, one must develop concepts of each of these various objects, kinds, properties and so forth. Developing adequate concepts of these kinds of things involves, paradigmatically, learning to reidentify them via their perceptual manifestations in a variety of ways, under a variety of conditions, given a variety of intervening media, and so forth. Because the proximal stimuli that may manifest the same distal object or property, and so forth, are extremely, numerous and diverse, depending on diverse outer conditions and intervening media, the business of producing new beliefs that correspond by uniform rules to the world with any regularity would be completely impossible unless abilities to recognize each of these various distal objects and properties in a variety of ways, to suit different mediating conditions, were itself highly developed. The representation-making mechanisms must be very versatile in recognizing, or better, reidentifying, the same object, the same kind, the same property or relation, and so forth as the same one again under lots and lots of different external circumstances.

How can I have concepts of things I have only heard about and can't recognize by normal perceptual means? I will get to that below. But for a full exposition and defense of this view of what having empirical concepts entails, see (Millikan 1984), Chapter 15 and (Millikan 2000), chapters 1 through 5.

The question, how the developing intellect learns to make these reidentifications correctly, concerns what can be called the natural epistemology of concepts, as distinguished from the natural epistemology of judgment or belief. It concerns the developmental processes by which children (and perhaps some animals) Normally acquire, through experience, an inner representational belief system in which the concepts or mental terms are non-equivocal, non-redundant, and non-empty, and a system which the belief-using devices have, at the same time, learned how or been tuned to use. That story is told in (Millikan 2000) chapter 7. It is a story that requires realist ontological underpinnings. These are provided in (Millikan 2000), chapter 2. But that there is a Normal way for children to acquire concepts or "mental terms" that are non-equivocal (no failures to recognize that Tweedledee and Tweedledum are two persons, not one), non-redundant (no Oedipal failures to identify Jocasta with Mom) and non-empty (no phlogiston concepts) does not have the least tendency to imply that no such faulty concepts are in fact employed by the average adult, nor that their employment implies that something was wrong with the minds of the persons that developed these faulty concepts. On this, compare (Millikan 2000) chapter 14.

What do we say, then, about mental terms that are defective? Well, if they are empty, if they do not exemplify any real ability to reidentify anything real, we say that they don't refer (counting real properties, real kinds and so forth as "referents" for ease of exposition). We don't say, for example, that they refer to possible entities, any more than we say that a can opener that isn't working correctly is opening possible cans. If the defective terms are equivocal, we can just say that they refer equivocally. If they are merely redundant, they refer and refer univocally, but they are not in accord with what

might be called an important "regulative ideal" for the cognitive systems. For a central job of the cognitive systems is to manage to learn to reidentify what is referentially the same thing as the same again.

Should we say that defective mental terms have "meanings," or that persons who harbor them mean something when use sentences that express them? They usually feel the same to the mind as mental terms that are not defective. They were developed by systems whose jobs were to develop (Normally functioning) mental terms. But there isn't any kind of entity with a rightful place in our ontology that corresponds to them. No possible individuals, possible kinds or possible properties; no Fregean senses; no Carnapian intensions; no Quinean meanings-under-glass-in-the-museum. Quine was right about all that. And Gareth Evans was right when he pointed out that the idea that there might be modes of presentation that didn't present anything made no sense at all.

But what do we say about, not faulty mental terms, but, granted there are such things, empty mental descriptions composed of healthy parts --say, the thought the present king of France? It certainly doesn't correspond to any peculiar entity such as a possible person. No entity without identity, and possible persons have no associated criteria of identity (Quine again). Nor does it correspond, of course, to an actual entity, to a referent. Each of its parts and its mental syntax are all perfectly fine. But if this is a real thought, and not just a pretend thought (such as the "thought of Santa Clause" for us adults), healthy parts or no, it is defective as a whole. It can, however, be expressed using healthy English words and healthy syntax. For this reason, the English phrase that expresses it can be translated. In that sense the phrase expressing it means something. Was there something more I was supposed to say?

Now, at last, I am in a position to try to mollify some internalist anxieties for the case of one kind of externalism, the kind that I have just outlined.

First, a useful analogy may help to make this kind of externalism more intuitive. No one doubts that whether the person you are apparently remembering at the moment is Great Aunt Nelly or Great Aunt Stella is a matter of the history of that particular memory, a matter of which aunt (if either) caused the apparent memory in that sort of way that memories are Normally caused. Similarly, the proposed kind of externalism claims that basic thoughts of things in the empirical world, such as thoughts of particular people and thoughts of water or of dogkind or of the color red, are rather like memories. What they are thoughts of depends on what in the outside world they were developed, in a Normal way, as a certain kind of cognitive response to. That is, what they are thoughts of, indeed, whether they are genuine thoughts of anything, rather than, say, just sorts of itches or tickles, depends on their causal histories.

In a way, this claim is just a broadened empiricism. The classical empiricists claimed that all thoughts have to have a certain kind of history involving the senses, and that this was not an accidental fact about human psychology, about what happens to cause the capacity to have ideas in humans, but a fact that followed from the very nature of what thought is. If a thought is a copy of an impression, for example, then it is a thought, in part, by virtue of its causal history. What I have changed here is that the relevant causal history is not the history merely of sensory input, but the history of the wider system in which the individual human head has been operating, a system that

has included aspects of its environment.

Now to this sort of externalist position, one standard internalist objection is that if historical events involving things outside your head determine what you are thinking of, so that what is presently inside your head leaves the matter indeterminate, it appears that you have no way of knowing what it is you are thinking of. To this objection, the first proper response is a question. What does the internalist suppose constitutes your knowing what you are thinking of?

Would you have to think of your thought, then think of its object, and then judge that the one represented the other? How would you think of the object in order to make this judgment? Would it just appear there before your mind, bare, the way Russell supposed sense data do? Or would it have some kind of essence that defined its real nature, which essence could appear directly before your mind? Or would you have to think it through another thought-representation of it?

Notice that the last option would seem to make the truth of your judgment depend on the fact, only, that two of your thoughts represented the same, rather than that either represented some particular thing having some particular defining nature. So long as the thought that you think of and the thought that you think with, in making the last sort of judgment, represented the same object, you would be judging correctly what you were thinking of.

Now what I have proposed is the following. There is no such thing as having an external object, kind or property directly before your mind. And I haven't the slightest idea what it would be to have something's nature before my mind, certainly without the thing, the kind, the property, or whatever, whose nature it was, being dragged along. True, an incredulous stare, as David Lewis was fond of pointing out, is not an argument. But I claim, positively, that something akin to the third alternative above is both readily understandable and plausible. To know what you are thinking of just IS to recognize when you are thinking of the same thing again.

Nor need this recognition consist in making a judgment. To recognize when your thought is of the same thing again is usually effected just by joining these thoughts, or being prepared to join them, in mediate practical or theoretical inference processes that use them as a middle term. These are processes that will perform in a Normal way to produce true conclusions or helpful desires or intentions only if these terms do in fact represent the same. For details on this proposal, see (Millikan 2000), Part Two.

That you know what you are thinking of when you think of an object, kind, or empirical property follows directly from the theory of empirical concepts sketched above. For on that theory, having a concept consists, in part, of the ability to reidentify the object of thought despite the variety of its manifestations to the senses, hence, of course, its possible diverse effects at the level of conceptual representation. A corollary is that one can know what one is thinking of more clearly or less clearly, depending on how versatile one is at reidentifying what one thinks of in perception and thought. Some of one's concepts are more adequate than others; some come closer to the ideal than others.

Nor does this conflict with common sense. We often know of a person under a certain description, such as "the president of the ACLU," yet are ready to say that we don't know who that person is. To know of a person under a certain description is, in

general, to know how to reidentify them in at least one way, namely should one encounter them again as fitting the same description. Indeed, usually a definite description affords one a variety of different ways to identify its referent, because one generally knows a variety of different ways to reidentify the correlate of each term in the description. But still, we may be ready to say we don't know who that person is. So what would be required in order really to know who that person is? On this question, I strongly recommend Böer and Lycan's book Knowing Who (1986), and I suggest that the lesson generalizes to all kinds of knowing what something is. One can know who one is referring to or thinking of in any of a variety of ways, the more the better, I suggest. The more ways one knows, the closer one is to the unobtainable ideal that would define perfect functioning of the one's concept-forming capacities.

Where the externalist theory I have presented parts company with common sense, however, is here. According to common sense, you know just by reflection whether you are thinking of something and if so of what. On the theory of empirical concepts I have proposed, whether you are thinking of something is tested by you over time empirically. That is what the naturalized epistemology of concepts referred to above is about (Millikan 2000, chapter Seven).

A second standard internalist's objection to externalist theories is that on such theories it is not possible to know with certainty what you believe. If what you believe depends on the history of your thought, not just on what is inside your head, how can you be sure what you believe? This may at first appear to be the same as the first internalist's objection just discussed. But the overlap is only partial. Knowing what you are thinking of, as I have just described it, is not the same as having beliefs about your beliefs.

Again, the first response to this objection should be a question, indeed, two questions. First, why is the internalist so sure that we always do know, with certainty, what we believe? Surely Freud rocked the complacency of even ordinary thinking on that point long ago. On the other hand, certainly it is true that people generally know in a lot more detail about what they themselves believe than about what anyone else believes. But people are also likely to know more about the details of what is currently in their mouths than anyone else knows. This is because they have a different way of finding out what is currently in their mouth than other people do. Similarly, it seems clear that we have a different way of finding out what we currently believe than other people do, and it seems to be a pretty accurate way. This brings up the second question we should ask the internalist.

By what sort of mechanism does the internalist suppose that we find out about our own beliefs? This ability could not be explained, of course, just by the fact that the beliefs were in our heads. Our neurotransmitters are in our heads too, but that doesn't make it any easier for us to know about them. It was, of course, traditional philosophical doctrine that the mind is somehow transparent to itself, that it does, somehow, know all about itself. All that is necessary to know about minds is to be a mind. But if that doctrine were true, why have we been disputing about the nature of thought and of mind for 2500 years? And why is it that it takes small children three or four years to learn to employ concepts of mental states, such as the concept of belief, appropriately? The internalist demands a story from the externalist about how we manage correctly to

recognize our own beliefs. But where has he given his own story on this?

Be that as it may, there was an interesting story on this told by the very first of the radical externalists, namely, Wilfrid Sellars. Sellars suggested that having acquired the ability to identify the presence of certain kinds of beliefs in other persons, one might later learn how to identify beliefs in oneself, roughly by the following means. One might catch oneself on the verge of candidly asserting something, and preface this asserting with either a spoken or a thought "I believe that...." To fill out the Sellarsian story, however, we would need to have a pretty definite view of the relation of thought to language. Sellars's own view was that the extensions of thoughts were determined by their inferential roles, including responses to external objects and properties with thoughts, and that inferential roles were internalized linguistic roles. That is very far from the position on thought that I have presented here. So to fill out Sellars's suggestion about how you find out what your beliefs are, I will need to fill in a different view of the relation of language to thought.

Earlier I mentioned thoughts of things you have heard about but would be unable to recognize in an ordinary way. How can you have concepts of these things on the theory of concepts I have suggested here? The proposal I have offered is that, Normally, language carries manifestations of distal affairs in the world in a way exactly analogous to the way light waves, sound waves, odors, tastes, gravitational fields and so forth do. Thus, in the Normal case (notice the capital 'N'), believing what you are told is in relevant respects exactly like believing what you see. Understanding and believing the messages carried by fact-stating language is merely another form of perception. One way of reidentifying an object or kind or property, then, is by knowing a word for it in the language spoken by the community you are in. And one way of having a concept, though not yet a very filled out or adequate one, may involve knowing just a word for its referent, rather than knowing what it looks like and/or feels like and so forth. Compare here the very slender concept one would have of the president of the ACLU if one had not the slightest idea who he or she is. Obviously this position on the relation of language to thought requires to be much clarified and to be argued for at some length. The clarification and arguments are in (Millikan 2000 Chapter 7).

Learning a public word for a thought referent is thus coming more fully to understand what you are thinking of when you think of that thing. Conversely, misunderstanding or using a public language word wrongly is way of failing to know exactly what you are thinking of.

Return now to the Sellarsian idea that you can come to know what you believe by catching your disposition to utter an informative sentence candidly, and prefacing the sentence with an overt or covert "I believe that....". Granted that you correctly use and understand the public language that you (propose to) speak, and that you are able to recognize your candid moods (this not to be taken for granted), you will express in language or understand in thought exactly what it is that you believe. No one else can come to know what you believe in this way. And we usually take it that this is a pretty accurate way of coming to know what you believe.

An interesting result follows. It follows that you could not know what you believed from inside if you had no language. That seems to me an entirely plausible result. Others may disagree. But in any event, if the internalist thinks there is some problem

that the externalist has but that he doesn't have concerning how a person comes to know what they believe, he must begin by giving us his own story on how this magic comes about.

References

- Böer, S.E. and W.G. Lycan 1986. Knowing Who (Cambridge MA: MIT Press).
- Block, N. 1986. "Advertisement for a semantics for psychology." In Studies in the philosophy of mind, Midwest studies in philosophy X, ed. P. French, T. Euhling and H. Wettstein. (Minneapolis: University of Minnesota Press).
- Dretske, F. 1986. "Misrepresentation." In Belief: Form, content, and function, ed. Radu Bogdan, 17-36. New York: Oxford.
- Dretske, F. 1988. Explaining Behavior. Bradford Books/MIT Press.
- Fumerton, R. 1989. "Russelling causal theories." In C.W. Savage and C.A. Anderson (eds.), Rereading Russell: Essays in Bertrand Russell's Metaphysics and Epistemology, Minnesota Studies in the Philosophy of Science, vol.12, 108-118. Minneapolis: University of Minnesota Press.
- Gallistel, C. R. 1990. The Organization of Learning (Cambridge MA: The MIT Press)
- Millikan, R. G. 1984 Language, Thought and Other Biological Categories. (Cambridge MA: MIT Press).
- Millikan, R. G. 1983 White Queen Psychology and Other Essays for Alice. (Cambridge MA: MIT Press).
- Millikan, R. G. 2000 On Clear and Confused Ideas; An Essay on Substance Concepts. (Cambridge: Cambridge University Press.)
- Millikan, R.G. 2002. "Biofunctions: Two paradigms." In Cummins, A. Ariew and M. Perlman, eds., Functions: New Readings in the Philosophy of Psychology and Biology, (Oxford: Oxford University Press).
- Millikan, R.G. forthcoming, Varieties of Meaning. (Cambridge MA: MIT Press).
- Neander, K. 1995. "Misrepresenting and malfunctioning." Philosophical Studies 79, 109-41.
- Putnam, H.. 1975. "The meaning of 'meaning'." In Minnesota Studies in the Philosophy of Science, vol. VII, ed. Keith Gunderson. Minneapolis: University of Minnesota Press. Also in Putnam (1975a).

-
1. Perhaps the best known example of this position is Block (1986).
 2. More careful definition and discussion of the notion "proper function" and "Normal" may be found in (Millikan 1984, Chapters 1 and 2; Millikan 1993 Chapters 1 and 2; Millikan forthcoming).
 3. Indeed, it is not merely an analogy if the position I have taken on the proper functions of human intentions hence of human artifacts is correct. See, perhaps best, (Millikan forthcoming, Chapter 1).
 4. The exposition I will give concerns, in fact, only one of three basic types of inner representation, namely "descriptive" or "indicative" representations. For more complete discussions see (Millikan 1984) especially chapter 6; (Millikan 1993) especially chapters 3 to six; (Millikan forthcoming) especially chapter 6 and chapter 11ff.

5. For more detailed formulations of this theory of inner representations, see (Millikan 1984), especially Chapter 6; (Millikan 1993), especially Chapters 3 to 6; (Millikan forthcoming), especially Chapter 6.

6. Millikan 2000, chapter 8 ff.