

# Perception, Action, and the Structure of Believing

John Perry

In R. E. Grandy and R. Warner (eds.)  
Philosophical Grounds of Rationality. (Oxford: Oxford University Press, 1986.)  
This version is from The Problem of the Essential Indexical and Other Essays,  
2nd edition (Stanford: CSLI Publications, 2000.)

October 22, 2011

Psychology has its beginning in ethological observation, Grice says. Systematic philosophical psychology appropriately enough has its beginning in ethological fiction, Grice's Toby and his fellow squarrels (1975). Toby and the other squarrels often gobble nuts in front of them, particularly after doing without nuts for some time. The theoretical apparatus we use to describe this in an explanatorily promising way allows us to say such things as *Toby prehends nuts as in front of him* and *Toby judges on nuts in front, for squarrel food*. Appeal to certain laws or near-laws then allows the explanation of Toby's gobbling. Prehending is a type of protoperception, and judging a type of protobelief. *Prehending nuts as in front of him* and *judging on nuts in front, for squarrel food* are states instantiated at different times, in front of different nuts, by different squarrels; gobbling is a type of act, similarly instantiated by different squarrels at different times; when successful, different nuts are gobbled. Thus the laws relate (proto) perceptual states, (proto) belief states, and types of action.

This is all, I think, just as it should be, not only for squarrels, prehension, judging, and gobbling, but also for humans, perception, belief, desire, and action. Psychological theory must deal with perceptual states, belief states, desire states, and types of action if it is to be general and systematic. The laws that relate these states and types of actions and their combinations to one another as typical causes and effects must square with ethological observation and make ecological sense.

A widely held conception of the structure of belief makes this impossible. The conception is the progeny of an inadequate semantics for attitude

reports due to Frege but attractive to many. This, in turn, is the result of an inadequate semantics for the sentences embedded in canonical attitude reports. The inadequate theory has been and is being developed in great complexity and detail, with impact on disciplines from metaphysics to syntax. When Grice's program of creature-construction reaches the point where it can give us a "genitorial" justification for the psychological laws or near-laws that govern us, we should have a reasonably clear idea of what these laws are. This requires, I think, a new semantic perspective.

I explain the problem and show how a new semantical perspective, situation semantics, promises a solution. Situation semantics, developed by Jon Barwise and me (see Barwise and Perry 1981, 1981a, and 1983),<sup>1</sup> brings back into semantics a notion banned by Frege (1892/1960, 64ff), the idea that sentences stand for something like facts.<sup>2</sup> It incorporates many of the insights of David Kaplan's three-tiered semantics, with its formal recognition of the sort of meaning that demonstratives and indexicals have (1978, 1979, and 1989). The importance of this type of expression has long been championed by a distinguished minority of philosophers of language, including C. S. Peirce, Arthur Burks, and Hector-Neri Castañeda (Burks 1949, Castañeda 1967). It is not surprising that a philosophy of language and mind that wishes to take our relation to the environment seriously must also take seriously *this*, *that*, *I*, *now*, and *here*.

This paper began as a discussion of the role of *I* in memory, occasioned by certain theses advanced in Grice 1941. A bit like the *monstera gigantea*, a vine mentioned later, the paper grew toward a nearby protuberance that cast a shadow on the original enterprise, leaving its roots behind. The protuberance was the semantics of indexicals and demonstratives. I end, however, with some remarks about *I* reminiscent of the themes that were once the core of the paper.

---

<sup>1</sup>The present paper was actually written in the fall of 1980, before the last two items mentioned, while actively working out the ideas of situation semantics with Barwise. It was heavily revised in the spring of 1982, incorporating some developments made in the theory in the meantime. The following papers, although not very up-to-date as far as situation semantics goes, explain the basic ideas and motivations reasonably well: Barwise 1981; Essays 1–4; Barwise and Perry 1980.

<sup>2</sup>Many authors have seen the need for facts or at least events in philosophy generally and semantics in particular. Our real situations are in many ways similar to Davidson's events, or at least are motivated by many of the same considerations and in response to Davidson's many insights concerning the need for events. Davidson was convinced by an argument we call the "slingshot" that events or situations could not serve as the semantic values of sentences, as they do in situation semantics; we discuss the argument in Barwise and Perry 1981. Our abstract situations are close to the Brandt/Goldman/Kim notion of event; see Goldman 1970 and Kim 1966.

## The Two Faces of Belief

Reports of the form *X believes that S*, where *X* designates a person and *S* is a sentence, appear to be of subject/object form. The subject is a believer, the object an entity designated by the that-clause, and the transitive verb is *believes*. We can speak of *X* and *Y* believing the same thing, of there being something that *X* believes, and so forth. The identity of the object believed or *proposition* seems to have much to do with the embedded sentence *S*.

Traditionally, propositions have been expected to play two roles. On the one hand, they identify states of the world. *It is true that Caesar was a Roman* makes a claim about the world, a claim that is either right or wrong depending on what the world is or was like. But propositions are also thought to identify belief states. Thus *Smith believes that Caesar was a Roman*, tells us something about the state of Smith's mind.

The dual role is connected with two uses we make of belief reports. We can use them as *evidence* about what the world is like: Turvey, a reliable authority on lunch, believes that lunch is served. So we believe it too, and go to lunch. We also use belief reports as parts of *explanations* about how people act: Turvey believes lunch is served; that is why he is leaving his office and running towards the cafeteria.

Propositions are identified by sentences, and thus seem to reflect logical relations among sentences (or perhaps impose them). So, using propositions to characterize the world, we can infer the truth of *that  $\psi$*  from the truth of *that  $\varphi$*  and *that  $\varphi$  only if  $\psi$* . But this logical structure is also useful in characterizing beliefs. I appeal to authority:

Our attitudes fit into a causal network. In combination, they cause much of our behavior; they are caused in part by the stimuli we receive from our surroundings and in part by one another. In attempting to systematize what we know about the causal roles of the attitudes, we find it necessary to refer to the logical relations among the objects of the attitudes (Lewis 1979, 514).<sup>3</sup>

---

<sup>3</sup>In this typically subtle and elegant paper, Lewis describes an alternative treatment of what I called a "self-locating belief" in Essay 1. Then, exploiting his modal realism, he treats all beliefs as self-locating. The modal realism does make a lot of streamlining possible. But I think our views on the structure of believing are very close. I suspect I do not agree with his remark in the footnote on page 541, that it is "unfortunate that the study of the objects of belief has become entangled with the semantic analysis of attributions of belief." This is supported by reference to the claims that "belief is in the head," and that "the main purpose of assigning objects of attitudes is . . . to characterize states of the head . . ." (526). Given the difference between belief as evidence about the world and belief

The phrases *objects of the attitudes* and *objects of belief* have three possible uses (at least). The ambiguity would be harmless, at least in the case of belief, if propositions really had the dual role envisaged. The phrase *object of belief* most naturally means those entities designated by the direct-object phrases in belief reports. So the object of Elwood's belief that Caesar was a Roman is whatever *that Caesar was a Roman* stands for. The phrase might mean those entities we are interested in when we use belief as evidence, the states the world either is or is not in. Or it might mean those entities that characterize our belief states in a way that allows us to systematize their causal roles, those entities in terms of which we should characterize states of the believer for purposes of explaining behavior. If propositions are the designata of that-clauses, identify states of the world, and identify states of believers, then propositions are the objects of belief in all three senses.

Experience with perception should raise doubts about this comfortable convergence of theoretical roles. Like belief reports, canonical perception reports identify a perceiver, a relation, and an object. But the point of the traditional arguments from illusion and perceptual relativity (properly construed) is just that the objects of perceptions in the first two senses are not the objects of perception in the third.

Suppose Smith says "Elwood sees Hoover Tower." *Hoover Tower* stands for Hoover Tower. That Elwood sees it might tell us something important about the world, most obviously that wherever Elwood is, Hoover Tower is there to be seen. But the facts of the relativity of perception show that there are countless ways to see Hoover Tower. The states we are in when we see Hoover Tower depend on our perspective and other conditions of perception, not just what we see. The actions that are appropriate for us, given certain desires, also vary with our position, relative to the object seen. So we should expect the ways of seeing Hoover Tower, and not the mere fact of seeing it, to be linked systematically with action and desire. A person who wants to reach Hoover Tower will turn one way if he sees it the way one does when one stands on the north side of it, and another way when he sees it as one does when one is to the south side of it.

Not only are there many perceptual states one could be in when seeing Hoover Tower, those same perceptual states could conceivably be involved in

---

as explanation of the believer's actions, these claims suggest an oversimple picture, which too easily leads us to think that our way of reporting beliefs is not "straightforward." Compare, "Vision is in the head." Surely, in some sense, true. The eyes and visual centers are all in the head. It does not follow that the point of saying what people see is mainly to describe their heads, or that, to the extent that we do report perception for that purpose, the expressions used will refer to what occurs in the head.

seeing something else. The perceptual state one is in when one sees Hoover Tower from one hundred yards on a foggy day in Palo Alto is quite a bit like that which one is in when one sees the Nebraska State Capitol Building from one hundred and fifty yards in Lincoln. We could use this fact to fool someone, were we so inclined and adequately funded.

*Objects perceived*, physical objects for the most part, do not serve well to classify perceivers for the purpose of explaining their behavior. This is recognized in commonsense psychology. We could never expect to explain what a person does just in terms of what they see, without, at least implicitly, considering the spot from which they see it. If I tell you that Smith wanted to get to Hoover Tower, and saw it from the Athletic Department, then I have given a reasonable explanation for his walking south. If he saw it from an airplane, I have not given a reasonable explanation for his walking south. It is also recognized that what is seen and the place from which it is seen sometimes give the wrong suggestion. If conditions are abnormal, one may not see things one usually does from that spot. In that case, we expect behavior appropriate to the spot from which things would ordinarily look that way.

In ordinary language, we have a variety of devices for describing our perceptual states independently of what we see when we are in them on a given occasion. For example, we explain mistakes by saying one thing looked like another; this can mean that we were in the state one is usually in when one perceives a thing of the latter sort. One can imagine a very systematic attempt to describe perceptual states, however. One might have a catalog of photographs, taken of various objects from various angles, and have people pick out one that shows “what things look like.” Such a system of identifying perceptual states, though less than the last word scientifically, might have its uses. Then the system of photographs used to individuate the perceptual states would be the “objects of perception” in the third sense indicated. This would be a very dangerous usage, however. For confusing the usages, but keeping the system of identifying perceptual states, one might begin to think of the photographs, or something akin to them, as what we *really* see. The third use of *objects of perception* seems best avoided altogether.

I think an approach similar to that described for perception (which is intended to be relatively noncontroversial), is needed in the philosophy of belief. The reason is that nothing can play both of the roles that propositions are supposed to play in the traditional theory.

The states an object is in at a time are distinguished from the relations in which it stands to other things at that time. That grass is green is a true proposition. It is not true just for certain persons or from certain positions,

it is simply true. Its being true, we might think of as a state of the world, rather than a relation that the world stands in to certain people, or certain positions within it. But if propositions correspond to such states of the world (and this is how I shall use the term), then there are no general and systematic links between belief in propositions, perception, and action.

To “systematize what we know about the causal role of attitudes,” we need states that are typically caused in normal perceivers by certain perceptual states and states that typically cause, in normal believers, certain kinds of actions. Belief in a proposition will not fill this theoretical need, so long as propositions identify states of the world. The propositions one comes to believe through perception will concern the objects in one’s environment; someone in some other place, looking at other objects, will not acquire belief in the same proposition, even if his perceptual state is the same. And the proposition, belief in which leads me to run or to reach out for a morsel of food, will not lead someone else, in some other place at some other time, to act in the same way.

Consider the proposition that the door to Building 90 at Stanford is open. Belief in this proposition can have no systematic and general links with action. Even holding desires constant for the general population, what one should do depends not just on the fact that the door is open and one’s desire (say) that it be shut, but where one is. Some with this desire and belief should walk downstairs and shut it. Others should phone and ask Julius to close it. A monk in Tibet, who has been apprised of the situation and desires that all doors be shut, really has no appropriate course of action open to him. Aristotle, who was awfully smart, may have foreseen that the door would be open, and have thought that philosophers should be kept warm. Apart from adding a footnote “Close the door!” to his *Ethics* (much studied in Building 90), it is hard to imagine what he could have done about it.

Standing in front of Building 90, one comes to believe that the door is open by seeing it open. But given the uniformity of Stanford’s Inner Quad, one would be in the same perceptual state when looking at Building 30’s open door (the numerals are obscured when the doors are open). But one would thereby be led to a different belief, that *that* door was open.

What we perceive—Building 90 or Building 30—is a matter not just of how we see things, but also of our position in the world. What we do—close the door of Building 90 or close the door of Building 30—is a matter not just of how we move but also of where we are and what we touch. If psychology, commonsense or otherwise, is to be systematic and general, it will be concerned with ways of perceiving and ways of acting, the same

theory applying to the philosophy major in front of Building 90 and the English major in front of Building 30. A similarly two-faceted account of the structure of belief is needed.

## Acceptance and Belief

Consider this unlikely conversation:

Sarah: It is Tuesday. So you should empty the trash.

Joe: I agree. I should empty the trash.

Jim: So you both believe that Joe should empty the trash.

I shall say that Sarah *accepts* the sentence *You should empty the trash*. By *sentence*, I mean the meaningful English sentence. Joe accepts *I should empty the trash*. Sarah does not accept *I should empty the trash*, but she thinks that when Sarah says it, she says something true. And Joe does not accept *You should empty the trash*, although he thinks that when Joe says it, he says something true. The sentences one accepts at a given time are those one uses or would use to describe the world to oneself and others. But description is just one activity among many; the sentences one accepts are those that guide one's actions generally. Some sentences, like *Caesar was a Roman* one usually accepts for a long time if one accepts them at all. Others, like *I should empty the trash*, one usually does not accept for long.

I suggest that the sentences one accepts give us a pretty good way of identifying the belief states one is in. I take belief states to be dispositions to do various things, including use various sentences in various ways, and so this is a case of identifying a multiply manifested disposition by one of its upshots. Thus belief states are identified by the behavior they cause in articulate adults in suitable circumstances. But belief states and belief are attributable to prelinguistic children and animals insofar as we are motivated to identify the states they are in with those of articulate adults. This is conceivable because the disposition to accept a sentence may be the disposition to do much else that is more important. This approach not only does not preclude the attribution of belief states to organisms without language, it facilitates it. Individuation of belief state by sentence accepted, rather than propositions believed, makes it possible to link some of these states systematically with perceptual situations on the one hand and environmental threats and opportunities on the other. This is precisely the step that is needed to make sense of belief without language.

Acceptance must be distinguished from belief. Belief is an attitude we have towards propositions in virtue of accepting sentences, or being in the

states that dispose us to accept them. Propositions I take to be abstract complexes of objects and properties; details are found in the next section. As Jim points out, Sarah and Joe believe the same thing: the proposition that Joe should empty the trash. The fact that they believe it in virtue of accepting different sentences manifesting different belief states, I shall put as they believe it *in different ways*. Of course, they both accept *Joe should empty the trash*, so there is a way they both believe it. But this is not necessary: one or the other might not know Joe's name. They might also accept the same sentence, and thereby believe different things. For example, if they both accepted *You should empty the trash*, there would be a sharp disagreement.

Acceptance is not belief that a sentence is true, although one will believe that one's uses of the sentences one accepts will be true, if one reflects on it. The analogy with perspective may be helpful. I have a certain perspective on Hoover Tower from my office. I realize that other positions give other perspectives, and I can imagine having those perspectives, and might even produce some rough drawings. I believe that those are how Hoover Tower looks from various angles. Of the drawing from the angle I have, I believe that it shows how Hoover Tower looks from where I am. And I can focus on the way it looks to me now, even without a drawing, and if I reflect upon it, I will certainly agree that this is how Hoover Tower looks from here, now. But believing that about how it looks should not be identified with its looking that way to me, now. Similarly, my acceptance of *Hoover Tower is over there*, should not be confused with my belief that the sentence, as uttered by one in my position and looking in the direction I am, is true. Accepting it is not thinking that it is true, but being in a certain kind of state identifiable by it.

When we describe beliefs, we usually identify the believer, the time at which the believing occurs, and the proposition believed. This is what Jim does in the example. A sincere first-person present-tense belief report—*I believe that such and such*—will embed in the that-clause a sentence accepted by the believer. But this is the exceptional case. Thus if Joe were to say “Sarah believes that I should empty the trash,” his belief report would be true, but he would not have embedded the sentence that Sarah accepts, which is not *I should empty the trash*, but *You should empty the trash*.

This may lead us to suppose that how a person believes a proposition is not very important. But often it is crucial. When we seek to explain a person's behavior by their beliefs it is ultimately the way they believe, not what they believe, that is important.

In the example, we may confidently expect Joe to collect the trash and head for the door, while Sarah continues her studies. Their behavior thus differs dramatically. But their beliefs are the same. It is the difference in



what they accept that appears to coincide with the difference in behavior. All well-behaved children will behave as Joe does when they accept *I should empty the trash*, even though each will believe something different.

Consider the English major and the philosophy major of the last section. They are in the same perceptual state, the one a normal perceiver is in, looking at a building of a certain type with an open door. But, of course, they see different things. One sees one open door, the other sees the other. Each is led to a belief state that is systematically and generally tied to such a perceptual state in a wide range of human beings, that state identified by acceptance of *That door is open*. They believe in the same way, though, of course, they believe different things. The one believes the one door is open, the other believes the other door is open. The belief state, given certain desires (where a similar distinction will have to be made), is linked with a certain type of action: closing that door. They both perform actions of this type, but, in so acting, close different doors.

The sentences whose acceptance will be systematically linked to perception and action typically contain demonstratives and indexicals: *That rock is coming at me*, *I'd better duck*, *This door is open*, etc. We need an account of meaning that comprehends such expressions.

## Situation Semantics

The framework I shall use is the situation-semantics theory that Jon Barwise and I are developing. The basic idea of situation semantics is that the meaning of a sentence is a relation between situations. First we shall look at the notion of a situation, then at the relation theory of meaning.

### Situations

The basic metaphysical idea is that reality consists of real situations, objects having properties and standing in relations at space/time locations. But we should not think of real situations as made up of separately existing objects, relations, and locations standing in some higher-order relation. Rather, objects, relations, and locations are abstractions from the flux of real situations, from what there ultimately is. They are at the most basic level of abstractions, but they are abstract. Such abstraction is the only way we have of dealing with the uniformities in reality, the recurring pattern of adjustment to which, whether through Divine Plan, evolution, or practical reasoning, is a precondition to effective action.

We assume that each situation  $\mathbf{s}$  has a location  $l$  ( $l = loc(\mathbf{s})$ ) and a type  $s$  ( $s = type(\mathbf{s})$ ). The type is a relation between  $n$ -place relations,  $n$  individuals, and 1 and 0 (whose role in this scheme will be explained shortly). Where we have a real situation  $\mathbf{s}$  with  $loc(\mathbf{s}) = l$  and  $type(\mathbf{s}) = s$ , then if

$$s(\text{runs}, \text{Albert}) = 1,$$

Albert is running at  $l$ . Situation types, like objects and properties and locations, are abstract; they represent *uniformities* across real situations. The type  $s$  just mentioned represents what all situations in which Albert is running have in common, while  $s'$  represents what all situations in which Albert is not running have in common, where

$$s(\text{runs}, \text{Albert}) = 0.$$

Now one can see the point of the 1s and 0s. They allow us to distinguish situation types that represent Albert as not running from those that are merely silent on the matter. Consider these two situation types:

$$\begin{aligned} s(\text{barks}, \text{Mollie}) &= 1 \\ s(\text{barks}, \text{Fido}) &= 0 \\ s'(\text{barks}, \text{Mollie}) &= 1. \end{aligned}$$

The first represents Mollie as barking and Fido as not barking; the second represents Molly as barking but does not represent Fido as doing anything.

A pair  $\langle l, s \rangle$  of a location and a situation is an *abstract* situation. An *actual situation* is an abstract situation that corresponds to a real situation. That is,

$$\langle l', s' \rangle \text{ is actual if there is a real situation } \mathbf{s} \text{ such that } loc(\mathbf{s}) = l' \text{ and } type(\mathbf{s}) = s'.^4$$

Other abstract situations are nonactual. Actual and nonactual situations are just set-theoretical objects that do or do not correspond to real situations. In particular, nonactual situations fail to correspond to real ones, rather than corresponding to nonreal ones.

There is, then, a gap between real situations, the stuff of the world, and abstract situations, actual and nonactual. The difference is obscured by calling them all *situations* and also by using, as I shall, the variable  $\mathbf{s}$  to range over both. *Situations* will usually be used for abstract situations, with *real* added when needed.

---

<sup>4</sup>This way of putting things builds in a decision about the question of how many situations there are at a location. If we think there is only one, whose type includes everything that is going on at a location, then we should want to require that  $s'$  be part of  $type(\mathbf{s})$ .

## Linguistic Meaning

In situations semantics, the linguistic meaning of a sentence is a relation between situations, between an utterance on the one hand, and a described situation on the other. Suppose someone says, “This table was built by you.” For this *utterance* to be true, a certain relation will have to hold between the utterance of it and another situation. Let us call the utterance  $\mathbf{u}$  and the situation it describes  $\mathbf{s}$ . Then the relation in question is this one:

There are objects  $a, b, c$ , locations  $l, l'$ , and types  $u$  and  $s$  such that:

- (i)  $\mathbf{u} = \langle l', u' \rangle$
- (ii)  $\mathbf{s} = \langle l, s \rangle$
- (iii)  $l$  temporally precedes  $l'$
- (iv)  $u(\text{speaks to}, a, b) = 1$
- (v)  $u(\text{demonstrates}, a, c) = 1$
- (vi)  $u(\text{table}, c) = 1$
- (vii)  $s(\text{builds}, b, c) = 1$ .

This is a complex relation that holds between abstract situations, involving many different individuals and locations. A true utterance occurs when there are situations in this relation that each correspond to reality: an actual utterance of *This table was built by you*, and an actual, properly related episode of table-building.

A theory of linguistic meaning assigns to expressions of a language relations between utterances and elements of described situations, in such a way that the correct relations for sentences emerge.<sup>5</sup> We indicate the meaning of an expression  $\alpha$  with  $\llbracket \alpha \rrbracket$ . This is a relation between utterances and situations if  $\alpha$  is a sentence, otherwise, elements of situations. Thus

$\mathbf{u} \llbracket \alpha \rrbracket \mathbf{s}$

is an instance of the form

---

<sup>5</sup>I am simplifying here, by sticking to situations, rather than bringing in *courses of events*. A course of events is a set of situations, hence a set of pairs of locations and situation types, hence a relation in extension between locations and situation types. Such a relation can be thought of as a partial function from locations to situation types. So a course of events is a dynamic version of a situation, it represents what is going on at locations, or, more precisely, part of what is going on at some locations.

$aRb$ .

Here are some semantic rules of this form,

- (a)  $\mathbf{u}[\textit{built by}]r$   
iff  
 $r$  is the relation of being built by.
- (b)  $\mathbf{u}[\textit{you}]b$   
iff  
there is an individual  $a$ , a location  $l$ , and a type  $u$  such that  
(i)  $\mathbf{u} = \langle l, u \rangle$   
(ii)  $u(\textit{speaker}, a) = 1$   
(iii)  $u(\textit{addresses}, a, b) = 1$ .

Note that “ $\mathbf{u}$ ” does not appear on the right-hand side of (a). *Built by* stands for a certain relation, independently of the facts of the utterance.<sup>6</sup> But *you* works very differently, since whom a use of *you* stands for depends on whom the speaker is addressing.

Here are the rest of the rules we need to handle our simple sentence.<sup>7</sup>

- (c)  $\mathbf{u}[\textit{table}]p$   
iff  
 $p$  is the property of being a table.
- (d) Where  $\alpha$  is a common noun, *this*  $\alpha$  is a noun phrase;  
 $\mathbf{u}[\textit{this } \alpha]c$   
iff  
there is a location  $l$ , a type  $u$ , a property  $p$ , and an individual  $a$  such that  
(i)  $\mathbf{u} = \langle l, u \rangle$   
(ii)  $u(\textit{speaker}, a) = 1$   
(iii)  $u(\textit{demonstrates}, a, c) = 1$   
(iv)  $\mathbf{u}[\alpha]p$   
(v)  $u(p, a) = 1$ .
- (e) Where  $\alpha$  is a transitive verb and  $\beta$  is a noun phrase,  $\alpha\beta$  is a verb phrase;

---

<sup>6</sup>Although we could treat the language spoken as an additional parameter of the utterance, and for certain purposes this is useful.

<sup>7</sup>Of course, these rules are not intended to be the final word, just to give the flavor of what a theory of meaning looks like in situation semantics. Rule (f), for example, would not fare very well as the fragment got larger.

$\mathbf{u}[\alpha\beta]a, \mathbf{s}, l$

iff

there is a relation  $r$ , an individual  $c$ , and a type  $s$  such that

(i)  $\mathbf{u}[\alpha]r, \mathbf{u}[\beta]c$

(ii)  $\mathbf{s} = \langle s, l \rangle$

(iii)  $s(r, a, c) = 1$ .

(f) Where  $\beta$  is a verb phrase, *was*  $\beta$  is a full verb phrase;

$\mathbf{u}[\textit{was } \beta]a, \mathbf{s}'$

iff

there are locations  $l, l'$ , and a type  $u$  such that

- (i)  $\mathbf{u} = \langle l, u \rangle$ ,
  - (ii)  $l'$  temporally precedes  $l$ , and
  - (iii)  $\mathbf{u}[[\beta]]a, \mathbf{s}', l'$ .
- (g) Where  $\alpha$  is a noun phrase and  $\beta$  is a full verb phrase,  $\alpha\beta$  is a sentence;  
 $\mathbf{u}[[\alpha\beta]]s$   
iff  
there is an individual  $a$  such that  
 $\mathbf{u}[[\alpha]]\mathbf{a}$  and  $\mathbf{u}[[\beta]]a, s$ .

### Meaning, Interpretation, and Information

Utterances are ways of conveying information, and although there are all sorts of information that an utterance can convey, the central case is that in which the utterance is true, and the information conveyed is that the conditions for the utterance's being true are met. On the relation theory of meaning, the truth conditions of an utterance pertain to two situations, the utterance and the subject-matter situation. They must be related in a certain way, and both be actual, for the utterance to be true. And there are cases in which the information we gain is limited to this. Suppose I get a postcard in the mail, with no signature, no return address, no picture to indicate where it came from, and the postmark blurred. Written on it are the words *I am having a good time*. What constraints are placed on the world if I assume that the writing of this card was a true utterance? Just this,

There are actual situations  $\mathbf{u}$  and  $\mathbf{s}$ , a type  $u$ , location  $l$ , and an individual  $a$  such that

- (i)  $\mathbf{u} \llbracket I \text{ am having a good time} \rrbracket s$
- (ii)  $\mathbf{u} = \langle l, u \rangle$
- (iii)  $u(\text{writes, this card, } a) = 1$
- (iv)  $l$  temporally precedes the present moment.

By knowing the conditions under which an utterance is true, and knowing something about the utterance, we learn about the subject matter, the described situation. Thus, if I know that you wrote the card last week, I learn that you were having a good time last week. This is clearly the normal pattern, knowing about the utterance, and learning about the subject matter. By fixing the facts of the utterance, and only allowing the situation to vary, we get the notion of interpretation:

The interpretation of an utterance  $\mathbf{u}$  of an expression  $\alpha$ , is the set  $\{\mathbf{s} \mid \mathbf{u}[\alpha]\mathbf{s}\}$ .

Consider my utterance to Jane, of *This table was built by you*. There are many situations in the interpretation of this utterance. The utterance is true if any of them are actual. They all have Jane building this table at some time previous to my utterance, but they vary in every other conceivable way. Some have World War II avoided by timely diplomacy, others have it happening as it did, others do not consider it.<sup>8</sup> Each of these situations has a type that is defined upon Jane and this table; these objects are the subject matter of the utterance. Nothing else makes it into the subject matter, not even me.

This last point is important. When we use natural language, we tend to focus on interpretation. The most common notion of saying the same thing is just uttering something with the same interpretation. If you say, pointing at Jane, and picking up a reference to the table, *She built that table*, we would ordinarily say that you said just what I did, that Jane built the table. The meaning of the sentences we use is not the same. And if it were the same, we might not have said the same thing. If you were talking to Albert and used the sentence *That table was built by you*, you would not have said what I did, but something quite different. Even in reporting utterances, it is the interpretation we worry about, not the meaning. You could report what I said as *He said that Jane built that table*, or even *He said that Jane built the table in the living room by the fireplace*. You are not reporting the meaning of the sentence I used, but the interpretation of my utterance.<sup>9</sup> This point carries over into the other propositional attitudes, such as belief and knowledge. Given this focus on interpretation, it is natural to capture one thing that has been meant by *propositions* with the interpretation of the utterance of a declarative sentence, a set (or class) of situations. This is how I shall use the term in this essay. So used, it conforms to the first two uses of *object of belief*, propositions are the designata ( $\approx$  interpretation) of that-clauses, and are the entities we are interested in when we use belief reports as evidence about the world. But the propositions in this construal do not give us the fine-grained way of classifying belief states that we would need to systematize their causal role. That job, as we shall see, is played by meanings.

---

<sup>8</sup>Here the lack of courses of events in this exposition makes accurate statement of the point impossible.

<sup>9</sup>This thesis is defended at length in Barwise and Perry 1981a and at greater length, with changes, in Barwise and Perry 1983.

This focus on interpretation in natural language leads one into thinking that all of the information we get from an utterance is information about the subject matter. But this is just not so. When we hear or read sentences that we take to be making true statements, we must simultaneously build up a picture of the utterance and of the subject matter. Often what we learn about the utterance, the general situation of the speaker, is more important than what the speaker says. A hostess hears a child say “This chair is funny looking.” She knows a little about the utterance situation from the fact that she can hear the voice. Taking the statement as true, she can fix the interpretation of *this*; there is only one funny looking chair in the house. This, together with the meaning of *the chair*, fixes the utterance situation in more detail, and she shouts, “Do not sit on that chair, it is a valuable antique.” What she learned was not that a certain chair was funny looking, which she already knew, but that a child was near her favorite antique. Another example. Elwood’s brother has made it to San Francisco, but is lost. He calls and says, “This phone booth is next to a large tower that looks like a fire-hose nozzle.” Elwood learns where his brother is, though his brother has not *said* anything about himself at all.

If the normal focus on interpretations is not properly understood, the other ways of getting information can seem mysterious. Sentences like *I am Elwood Fritchey* and *This city is San Francisco* can seem rather odd. If Elwood says the first one, he has said something necessarily true. If anyone else says it, they have said something necessarily false. And when Elwood says the first, it seems to have the same interpretation as *Elwood Fritchey is Elwood Fritchey*, when anyone else says it. And yet when he says it, I learn something, for I learn who I am talking to.

Let us introduce the notion of the inverse interpretation:

The inverse interpretation of a use of  $\alpha$  relative to  $\mathbf{s}$  is that set  $\{\mathbf{u} \mid \mathbf{u} \llbracket \alpha \rrbracket \mathbf{s}\}$ .

Relative to any actual situation  $\mathbf{s}$ , the only utterances in the inverse interpretation of *I am Elwood Fritchey* are ones in which Elwood Fritchey is the speaker, and the only utterances in the inverse interpretation of *This city is San Francisco* are ones in which the speaker is demonstrating San Francisco. If you and I are at Powell and Geary, and you say “This is San Francisco,” not pointing to a place on a map, but just the area around you, I learn that we are in San Francisco.

If we equate meaning and interpretation, perhaps lumping them together under something called “truth conditions,” then information of the sort just noted will be hard to deal with. We may deny it any semantic status,



but this is a mistake because the constraints placed on the identity of the speaker and other aspects of the utterance situation are as closely related to the meaning of the sentence used as the constraints its truthful use places on the described situation. We may shove all the information we can get into the interpretation so that we think that what is said contains implicit reference to the speaker and so forth. These strategies are both instances of what we call the *fallacy of misplaced information*. Given a relational perspective on meaning, we do not need to misplace information. The truth of an utterance of a sentence with a given meaning puts a set of interrelated constraints on two situations, the utterance and the subject-matter situation. When we recognize a truthful utterance, or take one to be true that is not, we pick up information, or misinformation, about both situations. The usual case is to know a lot about the utterance and learn about the subject matter, but cases where the basic drift of information is in the other direction are quite common.

So far, we have taken one thing about the utterance, the expression, to be given. But this is not always so. Indeed, when you report my utterance of *This table was built by you* by saying “He said that Jane made that table,” your auditor learns the interpretation of my utterance without learning which expression I used. But the expression is constrained to be one that would have that interpretation, as used by me, in the situation I was in. Given further knowledge about that situation, you may be able to figure out just what words I uttered. This sort of reasoning, and its analogues with reports about what a person believes or knows, is very important, as we shall see.

## Perception, Acceptance, and Action

This account allows us to see why the acceptance of certain sentences can be systematically linked with perception and action. By a *context* let us mean a situation like an utterance, except for the speaking or writing of a sentence. Given the way we have been using *utterance*, there are still a lot of facts left: who and where the potential speaker is, and how he or she is connected to the wider world.<sup>10</sup> A context belongs to the inverse interpretation of

---

<sup>10</sup>The facts about an utterance that are relevant to interpretation can be divided in various useful ways. We have generally divided them into facts about the discourse situation and connections. The former are publically observable facts such as the identity of the speaker and the time and place of utterance. The latter are basically causal connections to objects in the wider world. The interpretation of a use of *I* requires only the first sort of facts, the discourse situation. The interpretation of a use of a proper name requires the

an expression relative to a situation just in case the utterance gotten by adding the productions of the expression would. Finally, let us say that a context that belongs to the inverse interpretation of an expression relative to some actual situation belongs to the *actual inverse interpretation* of the expression.

The acceptance of a sentence  $\varphi$  will be systematically and generally linkable to one's perceptual situation, if one's situation can often be determined to belong to the actual inverse interpretation of  $\varphi$  perceptually. This does not require that all situations in the inverse interpretation of  $\varphi$  be perceptually similar, but that there be a subset of those that are. Thus *This is a fir tree* can be directly linked with perception, even though some situations in its actual inverse interpretation could not be perceived to be in it because it is dark or the fir tree is disguised as an elm.

The beliefs that can be acquired directly through perception seem to be those that are about the perceptible properties of objects in one's environment. This can be so because there are sentences whose acceptance for each person guarantees a belief about just those objects in that person's environment, and whose acceptance therefore can be directly and systematically linked to the perceptual states one is in when one is in certain kinds of situations.

Let us distinguish between sentences with bland and rich inverse interpretations. Those with bland inverse interpretations are those that are, or come close to, being eternal sentences. Their utterance has the same interpretation, no matter who says them and where and what they are attending to. Their meanings are insensitive to the context. Such sentences cannot be linked with perception in a systematic way. This can only happen if the sentence has a rich inverse interpretation, if its meaning puts heavy restrictions on the kinds of situations in which it can be uttered truly.

Similarly, the acceptance of  $\varphi$  will be linkable with action if there is a type of action that is generally advisable for speakers in situations in its actual inverse interpretation. For example, every speaker in a situation in the actual inverse interpretation of *There is a rock coming at me* should duck. Again, we get at the truth of a simple idea, that beliefs about the threats posed and opportunities afforded by objects in one's environment

---

latter sort of fact, a lesson to be learned from the theories of Donnellan and Kripke. Some kinds of expression require both sorts of facts for their interpretation, but they still are relevant in ways worth keeping separate. For example, the object demonstrated may be constrained by facts about the discourse situation but further causal facts are probably also relevant; which of the various chairs in a room a speaker refers to with *this chair* may turn on which one the speaker is attending to.

lead directly to action.

From this perspective, we can see the source of the tension created by the notion of the “objects of the attitudes.” If they are to be true or false, corresponding to states the world might or might not be in, the common objects of belief of various people in various places and times, they correspond to interpretation of utterances. If they are to serve the purposes of psychological classification, they should correspond to meanings of sentences with rich interpretations. The class of those who accept as I do is not the class of those who believe as I do. Many believe that I have a deadline; many of them are fishing, attending plays, and the like. But most of those who accept *I have a deadline* sit in front of their typewriters, thinking hard, and feeling pressure.

From this point of view, the importance of sentences with bland interpretations may seem a bit of a mystery. This is because to see how the acceptance of sentences can be linked to perception and action, we have had to ignore, in thinking about language, the very thing that makes language so important to us, the ability it gives us to communicate and hold information in ways that are not tied in any very immediate way to perception and action. An animal or a child will duck if a ball is thrown at them; there seems little point in attributing an intervening belief state. The perceptual states can lead to ducking through the design of the organism without belief intervening. If the only sentences we had were those capable of being directly linked to perception and action, those sentences would be useless. Their having a use requires two things. There must be some natural dispositions to act in certain ways, given certain perceptual states;<sup>11</sup> the acceptance of sentences can be seen as taking over this causal role. Second, there must be some other way besides perception of the objects in one’s environment to come to accept these sentences. The acceptance of sentences puts another node in the network of psychological states, one that can have the effect of perception without all of its risks. It is this second factor that makes sentences with constant meanings so important.

---

<sup>11</sup>The term *state* suggests much too static a way of looking at things, particularly things having to do with perception; I use it in spite of that, because of another set of suggestions it has, where the states of a thing at a time are contrasted with a wider set of properties dependent on its relations to the wider world.

## On Interpretation

Sentences, or tokens of them, travel. They travel quickly through the air as sounds, and slowly through the mail as marks. They are published, stored away in libraries, and sometimes read, at later times and far-away places. The utterance is in general not the same as the situation of interpretation, the time and place when the utterance is understood, and information gained from it.

Let us say that to interpret a sentence heard or read or otherwise apprehended is to find a sentence with the same interpretation in one's own situation, as the apprehended sentence had in the utterance of origin.<sup>12</sup>

We can distinguish several kinds of interpreting.

*Interpreting up.* This is to find an interpreting sentence with a less sensitive meaning. My friend in San Francisco sends me a card on which she has written, "This city has dilapidated cable cars." I write in the draft of my travel guide: "San Francisco has dilapidated cable cars." Note that the sentence I find is not insensitive. It has tense and a proper name. But it is less sensitive than the sentence I read on the card; it has a constant or near-constant interpretation over a wider range of change in the context.

*Interpreting down.* This is to find a more sensitive sentence with the same interpretation. On a trip to San Francisco, I read in my Mobil Guide, "San Francisco has dilapidated cable cars." I write on my notepad, "This city has dilapidated cable cars." Or I think it. But I do not get on the cable cars I see.

*Lateral interpreting.* This is to find a sensitive sentence to interpret a sensitive sentence. My friend shouts, "You are about to be hit by a rock." I think, "I am about to be hit by a rock."

We can always interpret a sentence, if we know its meaning. I get the postcard described above with no signature, no postmark, and no picture. I interpret: *The author of this postcard was having a good time when he or she wrote it.* Such a sentence is pretty useless in forming expectations or

---

<sup>12</sup>This is a dangerous usage. Sentence-*types* do not have interpretations, but meanings; they have interpretations relative to contexts. But we can think of sentence-*tokens* as having interpretations in an extended sense, the interpretations of the utterance that produced them. Xerox machines and other things complicate the latter notion, but I shall ignore such complications. This may be the point to indicate that I try to italicize mentioned sentence-*types*, but use quotation marks with verbs like the *says* of direct discourse. There is a theory of indirect discourse, which is that it is the mention of a sentence-*type*, but I am not at all sure that theory is right. It must have something to it, however, or I would not have so much trouble adhering to the convention I have described.

guiding action. Theoretical semantics is the systematic development of a theory of such (practically) useless interpreting sentences, in some language previously explained, of an appropriately formidable appearance.

These three types of interpreting have different purposes, different advantages, different disadvantages, and different requirements, even within the general project of gathering information.

Upwards interpretation allows the retention of information through change of discourse situation. *San Francisco has dilapidated cable cars* will have the same interpretation as I move about the world and for others reading my guide book. *This city has dilapidated cable cars* does not have this virtue. Upwards interpreting requires relatively insensitive expressions for the subject matter and knowledge of the utterance.

Downwards interpreting is used to generate expectations and guide actions. It requires knowledge of one's own situation. I could go from accepting *San Francisco has dilapidated cable cars* to accepting *This city has dilapidated cable cars* because I accepted *This city is San Francisco*. Having done so, I expect to see such cable cars, and avoid getting into them.

Thus it is the very blandness of inverse interpretation that makes insensitive sentences unsuited for direct linkage to perception, that makes them eminently suited for storing information thus acquired.

These reflections lead to the following picture of the structure of accepted sentences, or *doxastic structure*, in a normal believer. It has three levels. At the top are the insensitive sentences. At the bottom are those suited for direct linkage with perception and action. Tying these two levels together are the orienting sentences like *This city is San Francisco* and *That person is Richard Sklar* and *I am Elwood Fritchey*. Such a structure may be said to be in *perfect equilibrium* when the structure is closed under strong logical consequence.<sup>13</sup> This happens only rarely, one assumes, but I use *equilibrium* loosely, to mean approaching perfect equilibrium. For example, if *That person is Richard Sklar* is accepted and so is *Richard Sklar helps tourists in distress*, then *That person helps tourists in distress* is accepted. A structure is properly oriented when the orienting sentences are true in the context, and more or less fully oriented when there are a good number of them.

Given these rather vague notions, we can list some vague expectations we have about doxastic structures: program stability and generality, relative fullness and stability of orientation, relative stability of belief.

---

<sup>13</sup> $\varphi$  is a *strong logical consequence* of  $\psi$  when the interpretation of  $\psi$  in any context is included in the interpretation of  $\varphi$  in the same context.

## Program Stability and Generality

The program of an individual relates perceptual states to doxastic or belief states, doxastic states to each other, and doxastic and affective states to types of expectation and action. This is our loose version of Lewis' vision, a picture that the Gricean program in philosophical psychology can help bring into focus.

A program is reliable if it takes us from perceptual states to the acceptance of sentences that are true in the contexts in which the speakers are likely to be in those perceptual states, from sentences that are true in contexts to others that are, and from accepted sentences to appropriate actions. I think of the notion of a reliable program as a generalization of David Kaplan's concept of something true in virtue of the logic of demonstratives. The latter is a sentence true in every context, even though it does not have the same true content ( $\approx$  interpretation) in every context. *I am speaking* is such a sentence. It is true, whoever says it, although what they say varies from speaker to speaker and time to time. A reliable program need not be so flawless, just something that works out in most contexts. And truth is not always the relevant measure of success. Consider *Never do today what can be put off until tomorrow*. It gives us different advice every day, usually good.

A more sophisticated notion of a reliable program would consider changes from state to state that are reliable as one's discourse situation changes in automatic ways, as time passes.

We expect such programs to be stable within individuals and across members of a species, and to make ecological sense. The first attribute makes the elucidation of such programs a reasonable part of empirical psychology, the latter accounts in part for the illumination provided by Grice's program.

## Relative Fullness and Propriety of Orientation

We expect people to recognize the objects in their environment, by and large, to know who they are and approximately where they are, and the like. This requires changes in the doxastic structure as time passes and they and other objects move about and change. If they are aware of the passage of time, and little else, there will be an orderly transition of the whole structure. If aware of their own movement, local changes occur at the top level where their own name occurs, and massive changes at the bottom two levels. If aware of changes in other objects (of a relatively noncatastrophic nature), local change occurs at all three levels. (Compare Gibson 1979.)

## Relative Stability of Beliefs

We expect beliefs by and large to remain the same, except as necessitated by changes in accord with the principles just adumbrated. Thus we expect the doxastic structure to change so as to preserve what is believed, the interpretation of the accepted sentences. But we do not expect generality of belief. We expect different individuals to have different beliefs, not mainly because of disagreements about what the world is like, but because different individuals interact with, and are concerned about, different parts of the world.

Commonsense psychology is a psychology of differences. It has locutions built to focus on the way in which individuals differ; this means it focuses on what is believed, and not just what is accepted, because by the first two principles it assumes that difference of belief gives rise to appropriately different action. Given the assumptions of normal programming and relative full and proper orientation, classification by what is believed, the interpretation of the accepted sentences, rather than the accepted sentences themselves, makes sense. Relative to these assumptions, what is believed isolates a deeper dispositional property than what is accepted does; the latter changes to accommodate the former. (See Essay 4 for more about this.) We also have a reasonably rich vocabulary for describing orientation or lack of it: *recognizes*, *knows who*, *where*, and the like. And we have ways of cancelling the presumption of orientation when disorientation occurs.

When we speak of orientation, we have in mind, first and foremost, orientation towards objects in the environment. This is the requirement for making effective use of beliefs held at the upper level of our doxastic structure. But it is a fact of human life that we have many ways of designating individuals and the modes of designation are tied to actions directly and indirectly through different parts of our doxastic structure. We can speak of orientation here too.

Elwood is at a party and wants to look up a passage in *Word and Object*. He stands next to Marcia. I say, “Elwood believes that Marcia has a copy of *Word and Object* in her backpack.” I know he believes this because I said to him “Marcia has a copy of the book you are looking for,” just before the party. You expect him to ask her for it, and are puzzled when he does not. But Elwood is not fully oriented; he does not know who Marcia is. He cannot downward interpret in an effective manner, and so he cannot use his belief effectively.

Later Elwood is at home. He wants to call the person whom he was first told about, and later came to learn was standing next to him, who knew

so much about Quine. But he does not know who that woman is; that is, he cannot interpret upwards in the way necessary to use the information, stored in an insensitive way in the phone book, which has a lot of Marcias in it. It has often been noted that *knowing-who* and like expressions seem to vary a lot in their applicability depending on context. But there is a system behind it. These expressions are used to describe an interpretive ability suitable for the task at hand.

Knowing-who is being able to interpret information effectively for the task at hand. So knowing-who is a species of knowing-how. As used with reference to humans, it is usually a knowing-how that is supported by acceptance of some sentence. But such interpretive ability is not always based on acceptance. We should add three more species of interpretation.

*Interpreting in:* Interpreting information made available in ordinary perception (as opposed to perception of sentences) into accepted sentences.

*Interpreting out:* using information in accepted sentences to act effectively in ways directed towards objects in one's environment.

*Interpreting through:* acting on objects as a direct result of perceiving objects.

The concept of orientation is applicable here also. A shark that perceives a flounder by detecting a change of ionization in the water through which he swims turns towards the flounder, gives chase, and eats it. A tree casts a shadow on a *monstera gigantea*. The vine grows towards the tree and then climbs it, leaving its roots behind. In both cases, it is natural to speak of orientation: information about environmental objects was converted by a reliable program into action, or something like it, directed towards the same object. These programs are very reliable, but by taking the shark or the vine out of its ecological niche we could produce something resembling disorientation; language is not necessary for identity problems, although it helps.

Neither the shark nor the vine need to interpret up. In each case, the constant presence of the object throughout the episode eliminates the need for anything like a constant expression. *Information that is always available in the environment need not be retained in the head.* Intentionality has its origins in the ecological description of organisms, the borrowing of terms for environmental objects to describe states of the organism. Such descriptions require systematic interaction with the environment, not internal representations of it.

If the tree towards which the vine grows is cut, the vine has to wait until a different tree casts a new shadow before it changes direction. The



shark we can imagine to be more sophisticated. After chasing one flounder, it heads back towards another, through whose perimeter of ionization it passed during the chase. To make sense of this, we have to attribute a little more to the shark, something a little like our doxastic structures, that enable it to continue to believe in a perceptually inaccessible flounder.

When the shark perceived the flounder, the information it picked up was as much about it as about the flounder; it is its direction and distance from the flounder that is crucial, not the absolute location of the prey. And in returning to the flounder of second choice, the shark had to keep track of where it had gone, not just where the flounder had stayed. Does this mean that the shark needs also to be credited with some primitive precursor of *I*?

## Some Remarks About *I*

The shark needs no precursor of *I*, no self-referring perturbation of shark consciousness, no self-specifying blot in shark vision. It needs none because it has no access to information about itself except through perception, and no use for the information, except in action. The difficulty in grasping this point has less to do with the difficulty of imagining what it is like to be a shark, than with seeing clearly the connection between the meaning of *I* and its role in perception, cognition, and communication.

*I* stands for the person who uses it. This simple rule hardly seems to invest *I* with sufficient meaning to give it very special place in self-knowledge that philosophers have accorded it. There seems to be an immediacy and salience to knowledge that we formulate with *I*. Most famously, Descartes begins his climb out of the pit of doubt with *I think*, not *Descartes thinks*, or *The author of the Meditations thinks*. Hector-Neri Castañeda has enriched philosophy with many examples of importance (1966, 1968). A version of one: Ivan Tovar, heir to a famous fortune, does not claim it, in spite of reading accounts in the papers of the search for Ivan Tovar, the riches that await him, and so forth. Why? Because he has amnesia, and does not accept *I am Ivan Tovar*. He does accept *Ivan Tovar has many riches awaiting him*, but not *I have many riches awaiting me*. So he does not act.

Some philosophers have thought such importance shows that the rule cited cannot be the whole truth about *I*, perhaps not even an important part of it. Frege thought that *I* must have a special sense for each of us, a sense that determines that person as reference when he uses *I* (1918/1967). Anscombe suggests that *I* should not be thought of as a referring expression at all, for she feels that to suppose that it is leads us to take it not to refer to persons but to selves, a metaphysical nuisance (1975).

The special importance of *I* is easily accounted for in terms of our framework. Those who accept *I have many riches awaiting me* do not all believe the same thing, they may even disagree. But they all believe something that makes it reasonable for them to take steps to obtain their riches. Membership in this class of people projects onto behavior, more or less. But membership in the class of people who believe that Ivan Tovar has many riches awaiting him does not project onto obtaining behavior. So, Ivan's membership in the second class leaves him just like a lot of other people, while his membership in the first would lead us to expect him to seek his riches.

We can get a lot of knowledge about ourselves through perception. And knowledge about ourselves is quite relevant to action. The meaning of the word *I* makes it peculiarly appropriate for identifying a certain causal role. Acceptance of *I*-sentences, although determining different beliefs for each of us, plays a very similar role for all of us. That is why the first person is indispensable for philosophers who want to isolate an important class of doxastic states.

But indispensable as *I* is for philosophers in identifying such states, it is not necessary for someone to have mastered *I* to be in them. Children learn to use their own names before they master the intricacies of *I* and *you*. Katie may respond to *Katie get in here* and insist *Katie wants a cookie* long

before she masters the first person.

It is not surprising that use of one's own name precedes the use of *I*, or *can*. *Katie wants a cookie* is not in general an effective thing to say if one wants a cookie. But for Katie, it works well. If she gets in the habit of using it whenever she is cookie-deprived, it will work well. We can make the same point about an expression like *this city*. As long as Katie does not travel, sentences like *It is raining in Palo Alto* can be linked to her perceptual state. When she begins to travel, we will have to break this link.

Katie won't ever cease to be a Katie. But she may meet other kids with the same name. She may meet adults who do not know her name, and whose names she does not know. In these situations, *I* and *you* must be mastered.

When a perceivable entity *a* is constantly in *X*'s environment, *X* can link a sentence with an insensitive expression for *a* to her perceptual states. The constant expression can play the role of a demonstrative or indexical. But the opposite is also the case. An indexical or demonstrative phrase can play the role of a constant expression in the retaining of information in just such a case. The distinction between a reliable sentence—one that makes a truth whenever uttered, but not always the same one—and an insensitive sentence becomes blurred. *This planet is not a bad place to live, all things considered* has a very secure place in the doxastic structure of all but the pickiest nonastronauts. Does it belong at the top level of our doxastic structure or at the bottom?

*I* has a similarly blurred status, even for astronauts. Wherever I go, I am always there, standing right behind whatever is in front of me, the one with a head I cannot see, and a body that disappears under a mustache.

The fact that we can get by with our name and not *I* is based on the same fact as that we can get by with *I* and not our name. Both are needed for more complicated forms of life. I need my name to look up my own phone number. But I need *I* to engage in Cartesian doubt. Among the things I might doubt is whether my parents fooled me as to my name. They might have trained me to say *Thomas E. Dewey wants a cookie* when cookie-deprived, as a sort of misguided patriotic joke. It seems easier to worry about this than to worry about whether I got fooled or confused about the meaning of *I*, but if I wanted to worry about that, my name might come in handy. But suppose I wanted to worry about both at the same time. How could I think of myself?

Even a name is not necessary. The gannet folds its fragile wings gracefully to its sides just before it hits the water, when it dives for a fish. It does not fold its wings when it sees other gannets about to hit the water. It knows the difference between *its* hitting the water and *that bird's* hitting

the water. But it seems wrong to think of the gannet really receiving or using information that makes specific reference to itself at all. It has no use for *I*. Indeed, the gannet, a bird with rather a long protruding neck, usually does not even have a decent amount of itself in its own visual field. If we wanted to paint a picture of what the gannet sees, it would just be a picture of water, with the faint outline of a fish below the surface. And what seems true for the gannet seems true for us. We often see parts of our body. But we do not need to see parts of our body to know where we are in relation to the objects we see in the environment.

In just those cases in which either a name or a constant expression will do, because what is designated is always there, neither is really needed. The entity can be built into the semantics of the informational medium, without any perturbations in the medium, any name or indexical or disruption in the sensory field to designate that entity.

Consider the language I use with my dog: *Come, Go, Food, Sit, Stop that*. The dog interprets these; her actions and expectations vary (somewhat) with what I shout. Quite apart from whether this or anything like it could qualify the dog as penetrating the essence of language, the dog surely gets information from my verbal activity. We could take the meaning of each of these expressions to be a relation between an utterance of mine and situations involving my dog. This would be reasonable, in that the dog takes those to be commands for her. It would be misleading in that she has no apparatus for taking them in any other way.

We might say that my commands are preinterpreted for the dog. She does not need to interpret them up, down, or laterally. Perception is kind of like that. The information that we get at a certain spot in the world is information about objects in the neighborhood of that spot in a form suitable for the person in that spot. As long as this is the only source of information we have about ourselves, we need no way of designating ourselves, indexical or insensitive. Our entire perceptual and doxastic structure provides us with a way of believing about ourselves, without any expression for ourselves. As soon as we begin to get information about ourselves in other modes that need to be laterally or downwards interpreted to be effectively used, we will need to take some expressions as standing for us; to do so will be to interpret the information in a certain way. If Katie comes when we call *Katie*, she has it right. She still does not need *I*. She only really needs that to produce information about herself, not to interpret it. The first indexical she will take to designate her will probably be *you*, as said by her mother, standing

before her, posing some threat or affording some opportunity.<sup>14</sup>

---

<sup>14</sup>Jon Barwise and Michael Turvey gave me a lot of help on this. Various parts of this paper were included in talks at Stanford University, Harvard University, the University of Washington, and the University of British Columbia. I am grateful to the philosophers of those places for a number of probing questions and helpful comments. This paper was first completed while I was at the Center for Advanced Study in the Behavioral Sciences. I am grateful for financial support from the center, the National Endowment for the Humanities, the Andrew Mellon Foundation, and Stanford University. Revisions have been influenced by comments from Michael Bratman, Hector-Neri Castañeda, Julius Moravcsik, Howard Wettstein, and others, including someone who returned a copy of an earlier version with very useful comments, but no name.