

Individuals in Informational and Intentional Content

John Perry

1990

[In *Information, Semantics and Epistemology*, edited by Enrique Villanueva. Oxford: Basil Blackwell, 1990, 172–189. Reprinted in *The Problem of the Essential Indexical and Other Essays*. New York: Oxford University Press, 1993.

Mont Blanc with its snowfields is not itself a component part of the thought that Mont Blanc is more than 4000 metres high.

–Frege to Russell, November 13, 1904

I believe that in spite of all its snowfields Mont Blanc itself is a component part of what is actually asserted in the proposition, Mont Blanc is more than 4000 metres high.

–Russell to Frege, December 12, 1904¹

In this paper, I shall defend Russell’s view that Mont Blanc, with all of its snowfields, is a “component part” or constituent of what is actually asserted when one utters “Mont Blanc is more than 4000 meters high,” and of what one believes, when one believes that Mont Blanc is 4000 meters high. I also claim, however, that a proposition that does *not* have Mont Blanc as a constituent plays an important role in the assertion and the belief that Mont Blanc is more than 4000 meters high.

Taken somewhat out of context, the quotes from Frege and Russell express insights that pull in different directions in the contemporary philosophy of language. Behind Frege’s remark is the insight that reference is not direct but mediated. When we think and talk about things, our thoughts and words are not in *direct* contact with the things thought about and talked about. The meanings of our words, and the cognitive roles of our ideas, do not in and of themselves determine their reference. They get at objects via some aspects of

¹Both quotations are from Frege 1980 (163, 169).

the objects; the words directly provide only a *mode of presentation*, an aspect of the object to which they refer. Thus reference depends in part on the properties objects happen to have. Our thought and words could be just as they are, and refer to other things, given different facts about the world.

To individuate the thought in terms of the object, passing by the mode of presentation, seems to be to misrepresent its nature and distort its causal role. It makes different thoughts, involving different modes of presentation, come out the same. And it suggests that when a mode of presentation fails to pick out an object, the thought of which it is a part does not exist. The perceptions that lead to a belief, and the actions that flow from a belief, causally mesh with ways of thinking about objects, not the objects themselves.

And yet Russell's remark expresses another insight. We do, in fact, typically individuate the propositions we assert and believe by the objects referred to, not the modes of reference. The evidence assembled for this doctrine of "direct reference"—some of which is sketched below—is extremely persuasive. (See Donnellan 1970, Kaplan 1978, Kripke 1980, Wettstein 1984.)² But must we then abandon Frege's point of view?

I argue that the tension between these insights is apparent and not real. In the first place, theories of "direct reference" are misnamed. Kaplan's account of indexicals, Donnellan's account of proper names, and Wettstein's account of demonstratives all see reference as mediated. In each theory, a use of a term of a certain kind refers to an object, because that object plays a certain role relative to that use. On Donnellan's account, a use of a proper name refers to the object that plays a certain role in the historical account of the use. In Wettstein's account of demonstratives, the referent is the object that is salient in a certain way to speaker and audience. In each case, the referent must have a certain property: it must play a certain historical role in the production of the utterance (Donnellan), or be the salient object (Wettstein), or stand in a certain relation to the context of the utterance (Kaplan).³

It is important to note that in each of these theories the mode of presentation is *relational*. The object referred to is the one that stands in a certain relation to the use of the word in a specific utterance.⁴ The referent of a use of "I" is the person who is the speaker of the utterance. The referent of the word "this" is the object that is salient to the audience of the utterance.

I believe the tension between the two insights disappears when we distinguish the *truth conditions* from the *incremental content* of assertions and beliefs,

²In my discussions of these theories below I oversimplify; in particular, the simple theory of demonstrative reference in terms of salience greatly oversimplifies Wettstein's increasingly subtle thinking on the complex issue of demonstratives.

³Kaplan has emphasized that he did not choose the term "direct reference" to indicate that reference was unmediated, but to indicate that the object referred to, rather than the mode of presentation that determines it, is involved in the proposition expressed by the statement of which the term is a part. See Kaplan 1989a.

⁴In the sense of a certain uttering; a specific event in which a speaker uses the words in question.

and distinguish what is central from what is peripheral in Frege’s insight. In reporting assertions and beliefs in the normal way, we focus on incremental content. The incremental content of assertions and beliefs often has individuals as constituents in a clear and intuitive sense, thus vindicating Russell. But to have adequate theoretical accounts of these phenomena, we must also deal with the truth conditions. The truth conditions of an assertion or belief will not have any remote individual like Mont Blanc as a constituent, but only a “mode of presentation” that such an object satisfies. It is thus more Fregean than Russellian. But any reasonable theory of “direct reference” will provide such modes of presentation. The referent will be the object that stands in a certain relation to the utterance itself. A coherent Russellian must be a bit of a Fregean.⁵ Frege’s own theory ruled out such relational modes of presentation. But this was not a consequence of his insight about modes of presentation, but of other elegant but implausible features of his theory of content.

Propositions

I shall speak of R(eferential)-propositions and A(tributive)-propositions. An R-proposition is one that has an individual as a constituent—what Kaplan calls a “singular proposition.” Conceived as an R-proposition, the proposition that Mont Blanc is more than 4000 meters high might be modeled by a pair consisting of the mountain itself and the property of being more than 4000 meters high. More generally, a simple R-proposition will consist of a sequence of an n -ary relation followed by n objects. An R-proposition $\langle Pxy; a, b \rangle$ is true iff the objects a, b, \dots stand in the relation $P(x, y)$. The objects and the relation are the *constituents* of the proposition.

Conceived as an A-proposition, the proposition that Mont Blanc is more than 4000 meters high will be modeled by a pair of properties. The first property is being more than 4000 meters high. The second would be some mode of presentation P' , the property of uniquely satisfying some condition, which in fact Mont Blanc uniquely satisfies. One might suppose this property is associated, in one way or another, with the speaker’s use of the term “Mont Blanc.” Perhaps it is a property assigned to “Mont Blanc,” by the speaker’s “idiolect.” Perhaps it is the weighted sum of the properties the speaker would use to identify Mont Blanc. Perhaps it is the property of being the object that plays a certain causal role in the history of the production of the token of “Mont Blanc.” Wherever P' may come from, the A-proposition $\langle P; P' \rangle$ is true iff something coinstantiates P and P' .⁶

⁵Of course, Russell came to be a bit of a Fregean, in the sense that he allowed only individuals with whom one was acquainted to be constituents of the propositions one believed, and became very restrictive about acquaintance.

⁶For symmetry, we should define an A-proposition as consisting of a sequence of properties and a relation, true if some sequence of objects that instantiate the properties, instantiates the relation. But this complication is not needed in what follows.

Properties can also have individuals as constituents. Such properties are sometimes called relational as opposed to purely qualitative properties. Being the fountain immediately north of Hoover Tower, for example, is a relational property possessed by Tanner Fountain. As I am using the term, an A-proposition need not be purely qualitative. The properties that are the constituents of an A-proposition may be relational. Consider, for example, the property of being Frege's favorite mountain. This is a mode of presentation, and the proposition $Q' = \langle x \text{ is more than 4000 meters high; } x \text{ is Frege's favorite mountain} \rangle$ is an A-proposition.

Now consider these three propositions:

$Q = \langle x \text{'s favorite mountain is more than 4000 meters high; Frege} \rangle$.

$Q' = \langle x \text{ is more than 4000 meters high; } x \text{ is Frege's favorite mountain} \rangle$.

$Q'' = \langle x \text{ is more than 4000 meters high; Mont Blanc} \rangle$.

Q is an R-proposition with Frege as a constituent. It is true iff he had the property of having, as a favorite mountain, one that is more than 4000 meters high. Q' is an A-proposition, which is true if the properties of being Frege's favorite mountain and being more than 4000 meters high are coinstantiated. It seems that Q and Q' have equivalent truth conditions, even though Q is an R-proposition and Q' is an A-proposition. I shall say that Q' is the result of *refocusing* Q . Assuming that Mont Blanc was Frege's favorite mountain, Q' and Q'' also bear an important relation to one another. I shall say that Q'' results from *loading* Q' , given the fact that Mont Blanc was Frege's favorite mountain. This is not a very elaborate theory of propositions, but it will do for our discussion.

The terms "referential" and "attributive" come from Keith Donnellan's distinction between two uses of definite descriptions. One way of taking Donnellan's distinction is that statements containing referential uses express R-propositions, whereas statements containing attributive uses express A-propositions (Donnellan 1966).⁷ The basic question is whether the properties used in the description are part of "what is said." Suppose Smith says "The tallest mountain in the Alps is over 4000 meters high," and that in fact Mont Blanc is the tallest mountain in the Alps. It is clear that the truth of Smith's utterance requires that there be a tallest mountain in the Alps, and that it be over 4000 meters high. This is so whether we think that Smith is expressing an A-proposition or an R-proposition. Taken either way, if his utterance is true, the definite description must be satisfied, and the object that satisfies it must be over 4000 meters high. So taken either way, the truth of Smith's utterance requires the

⁷Donnellan considers this interpretation of his distinction; it clearly does not incorporate everything he has in mind. I am setting aside one important aspect of Donnellan's distinction, having to do with cases in which the description does not really fit the person the speaker intends to express a proposition about. For an important discussion of Donnellan's distinction and a number of related issues, see Neale 1990.

truth of an A-proposition. But is *this* what Smith has said, the proposition he has expressed? Or has he expressed the R-proposition?

To answer this question, we engage in counterfactual thinking. Suppose the Matterhorn to be higher than Mont Blanc, and both to be more than 4000 meters high. In that case, Smith's utterance would have been true. But would he have expressed the same proposition he actually did express? Would the proposition he in fact expressed be true in this counterfactual situation? If we are taking Smith to have expressed the A-proposition, we will be inclined to answer "yes" to both questions. But if we are taking Smith to express the R-proposition about Mont Blanc, we will take the answers to be "no" to the first question, but "yes" to the second. In the altered circumstances, Smith would have referred to the Matterhorn rather than to Mont Blanc, and so would have expressed a different proposition. But the proposition he in fact expressed would still be true in the counterfactual situation, since Mont Blanc would still be more than 4000 meters high. The fact that it would not be the tallest mountain in the Alps is not relevant, since that is not part of the content of the R-proposition.

The results of such a counterfactual test are less clear in the case of descriptions than in that of names, demonstratives, and indexicals, where the distinction between the truth conditions and the proposition expressed is more clear cut. Suppose Smith says, "You are sitting," while addressing Jones, who is in fact sitting, and has been sitting for some time, and is quite unconcerned with Smith. For Smith's utterance u to be true, he must be addressing someone who is sitting. The A-proposition $\langle x$ is sitting; x is the addressee of u captures this truth condition. But this is not the proposition asserted. Smith has not *said* that he is talking to anyone, and Smith is not *saying* anything about his own utterance. What he said was just that Jones was sitting, something that was true before he began, and would have been true even if he had never spoken. What he said is not the A-proposition just mentioned, but the R-proposition $\langle x$ is sitting; Jones).

Now suppose Smith says, "Mont Blanc is more than 4000 meters high," the case Russell had in mind. Above, we imagined two modes of reference that might be involved. Neither seems to be part of what is said. Suppose the mode of presentation is the weighted sum of properties Smith would use to identify Mont Blanc. Smith has not asserted that Mont Blanc has these properties. Even if Mont Blanc does not have these properties, and Smith would never succeed in identifying it, it does not change what he has said. Suppose the mode of presentation is that property suggested by Donnellan's account of names, of playing a certain role in the history of the production of Smith's token of "Mont Blanc." Smith has certainly not said that Mont Blanc has this property; the proposition he expresses using this token might be true, even if the token had never been produced. So, it seems that Russell is right. It is the R-proposition that Smith expresses. Mont Blanc is a constituent or component part of what is actually asserted in the proposition "Mont Blanc is more than 4000 meters high."

But this view presents a number of problems. Although the notion of an R-proposition is coherent, the notion of asserting or believing an R-proposition can strike us as very odd. For consider an assertion or belief t that expresses the R-proposition $Q = \langle Px, a \rangle$. There must be something about t that makes Q the correct interpretation of it. What could that be? It is not hard to imagine the property P being connected to some aspect of an assertion or belief by some general rule—a convention of language, or a law about how our minds interact with properties in the environment. That is, an aspect or property of an utterance or belief is associated by a general principle with a property or aspect of something else. But how could such a general principle connect an aspect of mind or language with a specific individual? It seems that the connection would have to be mediated by some individuating property the individual has. And it seems that this individuating property is what is important, not the bare identity of the object. This intuition is very strong and has led to the view that there is something about direct reference that defies common sense. I believe things fall into place, however, if we adopt a certain view of meaning—a view I associate with Hume.

A Humean View of Meaning

Although in his *Inquiry* Hume says little directly on the topic of meaning, Hume's point of view lends itself to a simple theory (1748/1975). Hume is concerned with our habit of forming opinions about what will happen at one time, on the basis of perception of what is going on at other times. He sees only one way this might work: we apply general rules or constraints, connecting types of events. We observe these types to be conjoined at the observed times, and assume them to be conjoined at the unobserved times. When one sees an event of one type, one expects an event of the other. I see the ball left unsupported; I expect the ball to fall. I see the man eating bread; I expect he will be nourished. Note that the types are interdefined. It is the very ball I see unsupported that I expect to drop, the very man I see eating that I expect to be nourished thereby.

Hume recognizes a relation between types he calls “constant conjunction.” I take this relation to hold if, whenever there is an event of the first type, there is an event of the second type. Of course, no one is in a position to observe that types are constantly conjoined, since everyone's experience is limited. But observing a number of cases of conjunction leads to another relation between one's perceptions (impressions or ideas) of types. I will call this “mental conjunction.” When T and T' are mentally conjoined by an individual, that individual will expect an event of type T' after observing, or otherwise coming to believe in, an event of type T .⁸ Hume recognizes, and emphasizes, that an individual may

⁸Note that since the types are interdefined, the mental states will have to be also. That is, there has to be something about a perception of a ball being left unsupported and an

mentally conjoin types that are not constantly conjoined. To mentally conjoin T and T' is something like believing that T and T' are constantly conjoined. But it is more primitive than belief; it is not the result of a rational process, but simply of a habit that, if we are optimistic, we may assume nature has built into us because we live in a pretty orderly place where experience of past conjunction is a sign of constant conjunction. (Hume 1748/1975; see especially the last two paragraphs.)

In Barwise and Perry 1983, Jon Barwise and I speak of constraints, involving, and attunement. These notions can be understood as generalizations of Hume's point of view. A constraint is a state of affairs in which two types are related: one type of event involves another. Constant conjunction is required for involvement; a thoroughgoing Humean will take it to be *all* that is required, but we left this question open. Hume's theory of habit and custom is a simple theory of attunement. His theory allows attunement to constraints that are not facts; this is what happened to Russell's chicken, who did not expect its head to be chopped off on market day, because it had not happened before then.

Given these ideas, we can introduce the notion of what a fact indicates, relative to a constraint. Suppose we have a constraint that a thing of type T is also of type T' . Relative to that constraint, the fact that x is T indicates that x is T' . I will call these the indicating fact and indicated proposition. A given tree has one hundred rings. That indicates that it is at least one hundred years old. Hume is eating bread. That indicates that Hume will be nourished. The car is making an odd sound. That indicates that it will soon break down. Now consider the indicated propositions in each case: that a certain tree is one hundred and one years old, that Hume will be nourished, that my car will break down. Are these R-propositions or A-propositions? They are R-propositions, with the tree, Hume, and the car as constituents.

In these cases, however, the source of the objects in the R-propositions is straightforward. We are confronted with the indicating fact, in which an object has certain properties. Given various constraints, those properties tell us something further about that very object. The tree that has one hundred rings lived for at least one hundred years; the individual who is eating will be nourished; the car that is making noise will break down. In each case, the proposition captures the further properties indicated about a constituent of the indicating fact, relative to the constraints. There is no mystery about these cases of indication. The object that is a constituent of the R-proposition is a constituent of the indicating fact. The constraint does not have any specific objects as part of its content. But it gets applied to specific objects. It says that if specific objects have one property, they will have another.

expectation of a ball dropping that makes the latter an expectation that the same ball will drop. See Essay 4.

Pure and Incremental Informational Content⁹

Here is another example of indication. The veterinarian takes an x-ray of a certain dog—say, the Barwises’ dog Jackie. The x-ray indicates that Jackie has a broken leg. This is a more puzzling case. Jackie is not part of the x-ray. She is not part of the general principles that govern the interpretation of x-rays. So how does Jackie become constituent of the proposition indicated by the x-ray?

We are assuming that there are constraints that limit the possible causes for an x-ray being of a certain type, and further practices that can augment the constraints to get us the information we want. An x-ray of the sort in question, can be brought about only if rays fall in a certain pattern on it. Given the additional regularities that arise from the practices in the veterinarians office—the circumstances under which x-rays get taken—an x-ray can be exposed to just this pattern only if it was produced by the x-raying of such an animal. So, relative to the relevant constraints and practices, this x-ray being of the type it is, carries the information that it was produced by exposure to an animal with a broken leg: \langle there is a unique animal of which x was taken, and that animal had a broken leg; this x-ray \rangle .

We get from this proposition to the one about Jackie by refocusing and then loading. First we focus not on the x-ray itself, but the dog to which it was exposed: $\langle x$ had a broken leg; x is the unique animal of which this x-ray was taken \rangle . Then we load Jackie into the proposition, given that she was that very dog: $\langle x$ has a broken leg; Jackie \rangle . Why do we refocus and load in this way? Because it is Jackie we are interested in. We are ultimately interested in the dog, not the x-ray, and we know which dog is in question. We are interested in the x-ray up to the point at which we have determined what it shows about Jackie; then the x-ray, and Jackie’s relation to it, cease to be of interest.

The pattern here is extremely common. We examine one object—call it “the signal”—because of what it shows us about another, already identified, which is connected to it, more or less remotely, in some way or another. We are interested in what additional information the signal has to offer about this connected object. Given that Jackie was x-rayed, what does the x-ray show about her? The R-proposition associated with the signal is simply what it shows about the object whose connection with the signal led to our interest in it.

In these cases, it is natural to think in terms of three interconnected types. There is the type of the indicating fact—in this case that an x-ray exhibits a certain pattern. And there is the type involved in the indicated proposition—in this case that a dog has a broken leg. The third type is the connecting type. It is the type of a fact that connects what indicates and what is indicated. In this case, it is the fact that Jackie was the dog x-rayed. The constraint somehow relates three types:

⁹For more on the topics in this section, see Israel and Perry 1990 and Essay 14.

T : x exhibits such and such a pattern.

T' : x was an x-ray of y .

T'' : y has a broken leg.

We can fit these three types into our notion of a constraint in different ways. First, think about the case where you are going through old files and find an x-ray exhibiting the pattern in question. You are sure that this x-ray could only have been produced by exposure to an animal; no x-ray could just accidentally come to look like that. And given that it looks like that, the animal of which it was taken has a broken leg. Here the constraint is that T involves T' and T'' . If an x-ray looks like this, there was an animal it was taken of, who had a broken leg. By examining the x-ray, you have learned something additional about it: \langle there is a unique animal of which x was taken, and that animal had a broken leg; this x-ray \rangle . This is the pure information carried by the x-ray, relative to the constraint above. Now you might naturally refocus on the unknown dog rather than the x-ray: $\langle x$ had a broken leg; x is the unique animal of which this x-ray was taken \rangle .

We can also think of the constraint *incrementally*. Suppose we already know which dog the x-ray was exposed to. That is, we know the fact of type T' . What *additional* information does the x-ray's exhibiting such and such a pattern provide us with? Given that the x-ray was taken of Jackie, it seems that the x-ray's exhibiting such and such a pattern indicates that Jackie has a broken leg. It seems this does not even require as strong a constraint as the one above. Perhaps x-rays can accidentally look like that, but in this case you know it was taken of a certain animal, Jackie. You only need the constraint that T and T' involves T'' . Every x-ray and animal are such that, if the x-ray looks like this and it was taken of the dog, the dog had a broken leg. This seems like a natural way to look at it—but it is not using quite the notion of indication that we have explained.

Let us introduce the notion of a fact incrementally indicating that P , given another fact. (I use “indicates_C” for “involves, relative to the constraint C .”)

A fact f incrementally indicates_C that P , given f' iff for some T , T' , and T'' :

- (i) C is the constraint that T and T' involve T'' .
- (ii) f is of type T .
- (iii) f' is of type T' .
- (iv) P is the proposition that some event is of type T'' .

The notion of incremental indication is doubly relative. One fact carries incremental information *relative* to a constraint, *given* another fact. Thus, in our example, the fact that the x-ray exhibits such and such a pattern incrementally

indicates that Jackie has a broken leg, *given* that the x-ray was taken of Jackie, and *relative* to the constraint that if there is an x-ray of such and such a type, then a particular dog has a broken leg, given that the x-ray was taken of that dog.

Now, why would it be useful to think of the x-ray incrementally, rather than purely? Suppose the vet is holding Jackie. Jackie is in pain and the vet wants to help. She plans to put Jackie in a certain room. She knows who Jackie's owners are, and plans to talk to them. So she has access to Jackie via a number of "aspects" or "modes of presentation." The issue in which she is interested is whether Jackie has a broken leg. She knows that the x-ray being developed is of Jackie. But this "mode of presentation"—being the dog that was x-rayed—is relevant only in connecting Jackie with the x-ray. Aside from that it is not as important as the other facts she knows about Jackie in guiding her actions towards Jackie. The vet looks at the x-ray, and says "It shows that Jackie has a broken leg." She is describing the content of the x-ray, *given* that the x-ray was of Jackie, because the fact that the x-ray is of Jackie is only relevant in getting the property of having a broken leg associated with her. Once that has been done, that mode of presentation becomes essentially irrelevant.

On the other hand, in order to understand general principles about x-rays and how they work, one has to deal with their pure content. To learn how to take and interpret an x-ray, just about the most basic fact that one has to know is that the developed x-ray contains information about the animal of whom it was taken.

Truth Conditions and Incremental Linguistic Content

The Humean picture is very simple. What gets us beyond the confines of any matter of fact are general principles. Only relative to such principles do we have informational content. To consider the content of a fact, is to consider it in terms of some such principles.

Now think of the conventions of language, as associating types of utterances with the types of situations they describe, and thus as assigning truth conditions to utterances. These conventions also generate "content." This will not be informational content, but *intentional* content. Types of utterances are not constantly conjoined with the situations of the type that language conventionally associates with them. Sometimes people lie, and sometimes they just do not get things straight. But the mechanism of content is the same: general principles that take us from the type of one event to another type of event identified in terms of the objects involved in the first.

Thus abstracted and generalized, the Humean analysis of meaning can be applied to language. The constraints that are relevant to linguistic meaning are

not laws of nature, and are not taken to be such, even by the most credulous among us. When we hear an utterance, we do not habitually expect that the proposition expressed will in fact be the case. But we do habitually *entertain* this proposition. The source of the habit is not observation of regularities, but mastery of the conventions of a given language. We are in some sense attuned to them, for our train of thought follows the conventions, even if we do not have the concepts necessary to articulate them.

Suppose the veterinarian says, pointing to Jackie, “This dog has a broken leg.” Here the bearer of content is the particular utterance, the veterinarian’s production of a token. The content-giving principles are the semantical rules of language that tell us under what conditions utterances of various types are true. If we look at what the vet would ordinarily have been taken to say, it would be that Jackie had a broken leg. That is, if someone were to report the vet’s linguistic act by saying, “The vet said that Jackie had a broken leg,” that would be an accurate report. Can we analyze this case in terms of pure and incremental content?

Let us develop a fairly plausible although oversimple account of the syntactic, semantical, and pragmatic principles involved here. Here is our account of the sentence the veterinarian used:

- (1) Tokens of the form “This X ” are noun phrases. Tokens of the form “has a broken leg” are verb phrases.¹⁰
- (2) If a token is of the form XY , where X is a noun phrase and Y is a verb phrase, the token is a sentence.
- (3) An utterance of a token of “This X ” refers to an object y if y is the object of the kind designated by the utterance of “ X ” that is most salient to the audience of the producer of the token at the time of the utterance.
- (4) Every utterance of a token of “dog” designates the kind, Dog.
- (5) Every utterance of a token of “has a broken leg” designates the property of having a broken leg at the time of the utterance.
- (6) If a token t is a sentence consisting of a noun phrase followed by a verb phrase, an utterance of t is true iff there is a unique object designated by

¹⁰A couple of words about how I am using “token” and “utterance.” A use or utterance often involves the production of a token, and tokens are often only used once, by their producers. That is what I am usually assuming here, but it is an oversimplification that ignores important issues. Tokens can be reused for different utterances, even by different agents. In the epistemology of language, tokens are most commonly (more or less) directly perceived, while utterances provide most of the properties crucial for interpretation. To see how important the difference can be, think of receiving a letter in the mail. One perceives the token, but cannot arrive at its content without knowing a good bit about the circumstances of its production or use. The differences between written and spoken tokens, printing, xerography, word-processing, and electronic mail, all raise interesting issues about the relation between utterances and tokens.

the utterance of the noun phrase that has the property designated by the utterance of the verb phrase.

I shall use the term “condition of truth” for the condition that an utterance has to satisfy to be true. This is the semantical property associated by the conventions of English and the general principles of language with indicative uses of sentence tokens of a given form, independently of particular facts about the context. Given this usage, different utterances can have the same condition of truth—even utterances with different truth-values. According to the theory above, the condition of truth for any utterance u of “This dog has a broken leg” are that there be a dog, who is most salient to the audience of the speaker of u , and who has a broken leg. That is, when we fix the facts of language and apply the above rules, this is what is required for truth.¹¹

In contrast to “conditions of truth,” I shall use the term “pure truth conditions” for the R-proposition consisting of the utterance and its condition of truth. Thus, if C is the condition of truth for a token u , the proposition $\langle u; C \rangle$ is the pure truth condition of u . The truth conditions of an utterance u will be a proposition that is true iff the utterance is. All utterances of the same type need not have the same truth conditions, and, in particular, false and true utterances of the same type cannot have the same truth conditions. Given the theory above, we can distinguish between the pure, refocused, and incremental truth conditions of the vet’s utterance u :

Pure Truth Conditions of u :

\langle there is an individual that is the dog that is most salient to the speaker in x and her audience at the time of x and that individual has a broken leg; u \rangle .

Refocused Truth Conditions of u :

$\langle x$ has a broken leg; x is the unique dog most salient to the speaker of u and her audience at the time of the noun phrase in u \rangle .

¹¹Note that I am assuming that the nouns and verb phrases, plus the conventions of language, get us to kinds and properties, independently of contextual facts, while indexicals, demonstratives, and names do not get us to their referents without context. Hence the kind Dog and the property of having a broken leg are “loaded” into the condition of truth, but not the particular dog. This is surely an oversimplification, for there is a great deal of interminacy in our use of common nouns and verbs that must be resolved by context.

Incremental Truth Conditions of u , given that Jackie is the unique dog that is most salient to the vet and her audience at the time of u :

$\langle x$ has a broken leg; Jackie \rangle .

Our ordinary notion of “what is said” is the incremental truth condition. We do not think that when the vet says “This dog has a broken leg,” that she has *said* something about her own utterance, namely, that there is a dog that is most salient at the time it was produced with a broken leg. We think she has *said* that Jackie has a broken leg. But this does not show that the pure truth conditions are not as I have described them. It simply shows we must distinguish between the truth conditions of an utterance and the proposition expressed by an utterance.

The proposition we take to be expressed by an utterance is typically the incremental truth condition, *given* the facts that determine the reference of the indexicals, demonstratives, and names. Above, in giving the incremental truth conditions of u , we fixed the fact that fixes the referent of “this dog,” according to the simple semantical theory for demonstratives sketched in rule (3).

Though I have spoken of *the* incremental truth conditions, it is important to note that there is actually a *range* of incremental truth conditions, depending on which facts we fix. Consider an utterance u of “I like that dog.” If we fix the facts about the demonstrative and the indexical, we get a proposition of the sort $\langle x$ likes y ; a, b \rangle . But suppose we fix the facts about the demonstrative but not the fact about the indexical. Then we get a proposition of the form \langle the speaker of x likes y ; u, b \rangle . This is not the pure truth condition, and not the proposition expressed. And yet this proposition might be important, in getting at what a given listener believes on the basis of hearing the utterance. Suppose a person hears u in a situation in which they can tell which dog is being referred to, but cannot tell, without turning around and looking, who is speaking. Perhaps a particularly ugly and mean-looking dog is on a stage, salient to everyone. Someone produces a token, and our listener turns around to see who might have such strange taste in dogs. The proposition above might be important in characterizing the belief the listener acquires from hearing the token that motivates him to turn around. The point is that the “proposition expressed” is not the *only* sort of incremental truth condition that is relevant to the epistemology of language. If we assume that it is, then we shall suppose that theories of direct reference must have difficulties with the problems of “cognitive significance”—with the possibility of different cognitive attitudes towards utterances that express the same proposition. (See Wettstein 1986.) They need not. Theories of direct reference say that modes of presentation are not constituents of the proposition expressed. But they need not deny that modes of presentation are constituents of all sorts of other propositions, including the pure truth conditions that are relevant to the epistemology of language. Theories of “direct reference” will lead one to this denial, only if one assumes what I shall call “the principle of unique content.” This is the principle that the only proposition associated with

utterances by the semantics of a language is the proposition expressed. (See Essay 11.)

What, then, is special about the proposition expressed? The answer has several related parts. First, in the paradigm case of communication, the audience is in a position to use its own knowledge to “load” the referents of indexicals and demonstratives and names. In a paradigm case of communication, we know who is speaking and are in a position to see to which objects the uses of the demonstratives refer. Second, the paradigm intention in uttering a sentence like “This dog has a broken leg” is to secure relatively stable belief in a proposition of the form $\langle x$ has a broken leg; $a \rangle$, not in the propositions that constitute the pure or refocused truth conditions of an utterance. This is not to say that belief in, or some other positive cognitive attitude towards, these latter propositions might not be required to get a competent speaker of English to the belief in question. But that is not the goal. If the veterinarian utters “This dog has a broken leg” to her nurse, her intention is that the nurse treat Jackie as a dog with a broken leg throughout Jackie’s stay. The nurse will have to have a way of thinking about Jackie that allows her to recognize her in a variety of contexts. The mode of presentation of Jackie that the nurse uses in understanding the vet—being the unique dog salient to the vet and her audience at the time of the utterance—is not such a mode of presentation. Once the nurse has looked at Jackie and noted enough about her to recognize her, this mode of presentation will have done its job. The nurse need not remember exactly what property of Jackie’s the vet used to refer to her, when the nurse first learned of the broken leg. This particular mode of presentation drops away as irrelevant once it has “fixed the reference,” to use Kripke’s phrase. Third, it is the incremental truth conditions of an utterance, given the facts that determine the reference of its indexicals, demonstratives, and names, that are typically the focus of reports about what was said by the utterance. I am suggesting that the reason for this is not that this proposition is any more directly or intimately connected with the utterance than other candidates. It is less directly connected with the utterance than the pure truth conditions.¹² The reason is simply that our notion of “what is said” reflects concern with the paradigm communicative situation.

Once one abandons the principle of unique content, one comes to expect the notion of “the proposition expressed,” and “what is said” to have a certain looseness. Consider the example above, about the ugly dog. Suppose I have one friend who is well known to like ugly, mean dogs. He arrives at the dog show, fits into the crowd behind me, and says “I like that dog” as a way of announcing to me that he has arrived. He wants me to turn in the direction from which the token came, because I believe that the speaker of it likes ugly, mean dogs, and have inferred from that, that my friend with the unusual taste in dogs has arrived. He is not really trying to get me to believe that he likes that particular

¹²Note that one might take the fully incremental truth conditions of an utterance to be its truth-value. Truth-functional expressions would then focus on an even less direct, more incremental property of utterances than terms like “said.”

dog; expression of that proposition is just a by-product of his utterance. The proposition he is really trying to convey corresponds to the incremental truth conditions given the facts that fix demonstratives, but not given the facts that fix indexicals: \langle the speaker of x likes that dog; u \rangle . In such a case, is it so clear “what is said”? Or, consider the case in which someone comes up to me at a party and says, “I’m Elwood Fritchey.” I learn, and am intended to learn, that the name of the person addressing me is “Elwood Fritchey.” A theory of direct reference can handle this, once it has abandoned the principle of unique content, even if we take a proposition of the form $\langle x = y; a, a \rangle$ to be “what is said.” (See Barwise and Perry 1983, 167ff.) But perhaps we should go a step further and simply be a great deal more flexible about what is said. The question here is the status of an additional principle for our simple semantical theory:

- (7) The proposition expressed by an utterance u of a sentence consisting of a noun phrase followed by a verb phrase is its incremental truth conditions, given the reference of the noun phrase of u .

Is this part of the semantics of “This dog has a broken leg?” Or is it really just a part of the semantics of “says”? Or is it not even that, but simply the typical result of some more complex, pragmatically driven principles about “says”? I am not sure how to answer these questions, but I am leaning towards the last answer.

If we take this line, we will find many of Donnellan’s claims about the referential-attributive distinction congenial. We will see this as a distinction between cases in which the proposition expressed is to be incremental, given the facts determining the reference of the description in question, and cases in which it is not. Facts about the communicative situation of the particular utterance will dictate the answer, rather than the syntax or semantics of the particular utterance.

Truth Conditions and Incremental Content of Beliefs

I have argued in various places that we need a two-tiered theory of belief and other cognitive attitudes. (See Essays 1, 2, 4, 6; Barwise and Perry 1983.) We need to distinguish between the belief states and the propositions believed. The belief state a person is in has to do with the internal goings-on of the mind or brain. The proposition believed will depend also on various sorts of external circumstances.

Here is an example of this distinction. I think there is a certain mental state with the following typical causes and effects. It is typically caused by some combination of the following: feeling of pressure on one’s thighs, buttocks and back, visual perceptions of one’s body being supported by a chair, memories of

sitting down, and the like. It typically causes one to be disposed to answer “Yes” to the question, “Are you sitting?” in relatively normal circumstances. If I am in this state, I believe that I am sitting. If you are in this state, you believe that you are sitting. We are in the same belief state, but we believe different propositions.

Jerry Fodor has a theory of the nature of cognitive states that we can use to think about this distinction, even if we have some doubts that it is the final truth about such issues. On his theory, there are sentences in mentalese written in various parts of our brain. So, in particular, there is a sort of belief box, in which are various tokens written. Let us assume mentalese is a lot like English. Then, in the above example, you and I both have tokens of “I am sitting” written in our belief boxes. So we are in the same belief state. But we believe different things.

Let us call my belief token b_1 and yours b_2 . These two tokens have the same condition of truth:

A belief token b of “I am sitting” is true, iff the person, in whose belief box b is written, is sitting.

The two tokens have different pure and incremental truth conditions:

b_1 's pure truth condition is $\langle x$ is written in the belief box of a person who is sitting; $b_1 \rangle$.

b_1 's incremental truth condition is $\langle x$ is sitting; John Perry \rangle .

b_2 's pure truth condition is $\langle x$ is written in the belief box of a person who is sitting; $b_2 \rangle$.

b_2 's incremental truth condition is $\langle x$ is sitting; you \rangle .

What do we believe in this case? Each of us believes that he or she is sitting. I believe $\langle x$ is sitting; John Perry \rangle , you believe $\langle x$ is sitting; you \rangle . It seems quite clear that we do not believe the pure truth condition of our belief tokens. Most people who believe themselves to be sitting do not believe, of the token in virtue of which they believe themselves to be sitting, that it is in the belief box of someone who is sitting. And, even those rare folks who do, do not have *that* belief, in virtue of that token.

The following, possibly historical, example will bring out this point. One Monday, Jerry Fodor gave a lecture in a certain room at MIT. He believed his own theory and believed that he was standing. He dubbed the belief token of his, in virtue of which he believed he was standing, X . He wrote down various things on the board about X , including

X is a token of “I am standing.”

The pure truth conditions of X are $\langle x$ is a token in the belief box of someone who is standing; $X \rangle$. Call this proposition P . Now the next day, we may

imagine, Fodor happened to go into the same room and sit down, shortly after his colleague Ned Block has taught a class. In fact, Block had not used the board and it was still unerased from Fodor's last lecture. But Fodor assumed that what was written on the board was written there by Block. He assumed that Block had been doing just what Fodor did the day before, using a real-life example to illustrate the mental sentence theory. So Fodor believed that X was a token of Block's, and that Block was standing when he used it as an illustration. So Fodor believed that X was a token in the belief box of someone who was standing. Fodor now believes P , the pure truth conditions of his own token. My point is that he does not believe P in virtue of that token's being in his belief box. The belief a person has in virtue of having a token in their belief box is *not*, at least in the vast majority of cases, a belief about that very token.¹³

So, we find in the case of belief something parallel to what we found in the case of assertion. Just as what is asserted is not the pure but the incremental content of an utterance, what is believed is not the pure but the incremental content of the belief state. We are interested in what conditions a belief has to fulfill to be true, given the facts about whose belief it is, and whom the mental terms in the belief stand for.

But now consider the commonsense principles that we have, connecting beliefs, desires, and action. These will not be concerned with the proposition believed, but with the belief state, not with what a person believes, but with how they believe it. As in the case of informational and linguistic content, so too in the case of the content of the attitudes: in specific cases we are interested in incremental content, but in understanding general principles, it is pure content that is relevant.

On Being a Fregean

Now we are in a position to see the general strategy for accommodating both of the insights expressed by the quotes from Frege and Russell. Frege is right that a thought or an assertion about a thing requires a mode of presentation. Let R be the relation that, on Donnellan's theory, holds between the use of a name and the object to which it refers. If we accept Donnellan's account, the pure truth conditions of an assertion u of "Mont Blanc is more than 4000 meters high," will be $\langle x$ is such that there is a unique object that has R to the noun phrase of x , and it is more than 4000 meters high; $u \rangle$. The refocused truth conditions will be $\langle x$ is more than 4000 meters high; x is such that it is the unique object that has R to the noun phrase of $u \rangle$. The first constituent of this proposition is a mode of presentation of Mont Blanc. And some such mode of presentation is required. There is no *direct* way for a use of "Mont Blanc" to refer to Mont Blanc.

¹³This example is borrowed from Essay 14.

But this mode of presentation will not be a part of the proposition that the utterance expresses. Russell was right about that. It is the object referred to or thought about, and not the properties by which we refer to it or think about it, that is part of the incremental content, and it is the incremental content that is asserted or believed.

There are certain other Fregean principles that conflict with this way of accommodating things, however. Frege not only thought that modes of presentation were involved in speech and thought, he held that those modes of presentation were constituents of the propositions asserted and believed. That is, in the terminology I have been using, he thought that the proposition expressed by an utterance involving a term incorporated the mode of presentation associated with the term. And, in the case of belief, he thought that the proposition believed incorporated the modes of presentation involved in the belief. With some qualification, the accommodation I have described rejects that view. These modes of presentation find their place in the refocused truth conditions of an utterance or belief, not in the proposition expressed or believed. Frege assumed, quite naturally perhaps, that the content we would focus upon in reporting assertions and beliefs was that most intimately connected with the words used or the type of thought. But this is not how we typically report language and thought. We focus on incremental content, on what needs to be the case for the utterance or belief to be true, *given* various contextual facts. The qualification is that when we give up the principle of unique content, the whole question of the content of a statement or belief becomes a much less determinate question.

While Frege thought that modes of presentation were involved in the propositions expressed in speech and thought, it is clear that the “utterance reflexive” modes of presentation I have been discussing were not the ones he had in mind. Even if one held that such modes of presentation were constituents of the propositions statements express, it would not save all of Frege’s theory. Frege held that when one reports the belief or assertion of an agent, the proposition expressed by the content part of one’s report expresses the very proposition that is said to be believed or asserted by the agent. But suppose you say, “Mont Blanc is more than 4000 meters high,” and I report, “You said that Mont Blanc is more than 4000 meters high.” The refocused pure content of your statement, and the refocused pure content of the content part of my statement will not be the same, since the modes of presentation involved will have different utterances as constituents. This identity holds only at the level of incremental content, when modes of presentations have been left behind.

Frege’s theory seems to imply that the mode of presentation, by which we think about an object, is a part of what we believe about the object. This seems incorrect. I may be thinking and talking about George Bush, in virtue of the complex role George Bush plays in the history that led to my use of “George Bush” and my acquisition of the notion I use in thinking about George Bush. I need not have any belief that George Bush plays this role, in order to talk about him and believe things about him. One may have mastery of the use of proper names, without having the concepts necessary to articulate how they work.

But there is nevertheless an epistemologically central connection between the correct theory of how a class of terms work, and the relations that ordinary language users expect, between uses of those terms and the objects to which

they refer. In the case of indexicals, it is plausible to suppose that competent speakers of the language pretty much understand what this relation is, even if they cannot articulate it. A person who understands a use of “I,” faced with an utterance of “I am such and such,” will look for the speaker of the utterance to verify the remark. Even in the case of proper names, we expect some knowledge of that relationship. Consider, for example, someone who comes across a conversation in which the name “George” is being used. He is not sure who is being referred to—is it George Bush or George Washington or George the fellow down the street? A competent language user would have some idea what counts as evidence for and against each of these hypotheses. If the straightforward approach of asking were unavailable, he would begin to look at the historical factors leading up to the use of the term in the conversation. Were the participants watching a campaign speech? Visiting a Revolutionary War battlefield? The mode of presentation of a term, though it may not be part of the content of the beliefs of someone who understands a use of the term, will still be cognitively important. If someone has mastered language, they will be *attuned* to the principles that govern relations between uses of terms and their references, even if they cannot articulate them and do not have the concepts necessary to be said to believe them. Frege’s view, though incorrect, is not groundless.

A theory of reports of thought and language based on the ideas explored here, then, will not have the elegance of Frege’s theory. But the modifications we have to make in his ideas are not reason to abandon his central insight. Reference is not unmediated. There is nothing particularly direct about “direct reference.”¹⁴

¹⁴I am grateful to Leora Weitzman for very detailed and helpful comments on the penultimate version of this paper, and to Howard Wettstein, Syun Tutiya, David Israel, and Mark Crimmins for comments on earlier versions. The recent Stanford dissertation projects of Weitzman, Crimmins, Lisa Hall, Genoveva Marti, and Stephen Neale have influenced my thinking greatly on these matters. My work on this paper was supported in part by the Center for the Study of Language and Information at Stanford, through an award from the System Development Foundation.