Reference

1. The phenomenon of reference

In a paper evaluating animal communication systems, Hockett and Altmann (1968) presented a list of what they found to be the distinctive characteristics which, collectively, define what it is to be a human language. Among the characteristics is the phenomenon of "aboutness", that is, in using a human language we talk **about** things that are external to ourselves. This not only includes things that we find in our immediate environment, but also things that are **displaced** in time and space. For example, at this moment I can just as easily talk about Tahiti or the planet Pluto, neither of which are in my immediate environment nor ever have been, as I can about this telephone before me or the computer I am using at this moment. Temporal displacement is similar: it would seem I can as easily talk about Abraham Lincoln or Julius Caesar, neither a contemporary of mine, as I can of former president Bill Clinton, or my good friend John, who are contemporaries of mine. This notion of aboutness is, intuitively, lacking in some contrasting instances. For example, it is easy to think that animal communication systems lack this characteristic—that the mating call of the male cardinal may be caused by a certain biological urge, and may serve as a signal that attracts mates, but the call itself is (putatively) not **about** either of those things. Or, consider an example from human behavior. I hit my thumb with a hammer while attempting to drive in a nail. I say, "Ouch!" In so doing I am saying this because of the pain, and I am communicating to anyone within earshot that I am in pain, but the word *ouch* itself is not about the pain I feel. If, on the other hand I say, with unnatural calmness, "Pain is present in my thumb", then I am in this instance talking **about** pain.

Such intuitions have, for the most part, been extremely compelling, in fact so compelling that the CORRESPONDENCE THEORY OF MEANING has, since classical times, in one form or another, been by far the most persistently pursued notion of how meaning in language is best characterized. Not to put too fine an edge on it, this is quite simply the idea that the significance or import of natural language utterances is found in the ways

in which they correspond to facts and things in the world around us. In present times, this finds its clearest articulation in the framework of model-theoretic semantics. Yet not everyone finds these basic intuitions of aboutness quite so compelling as to base a theory of natural language meaning upon them. Most notably in this past century, Wittgenstein is generally interpreted as articulating quite a different view of natural language meaning which, at best, treats "aboutness" as derivative or epiphenomenal (Wittgenstein 1953). Also, Chomsky (1981, 1992, 1995), Hornstein (1984), Ludlow (2003) and others have articulated a similarly skeptical view about its centrality. Since this chapter is about (the notion of) reference, I set aside consideration of such alternatives and focus exclusively on work which does find this initial intuition most compelling.

The word *about(ness)* itself, however, is a folk notion that is too general and vague to really get at something fundamental about natural language. We may ask, quite sensibly, what is Beethoven's third symphony about, what is the relationship of a couple really about, what is a painting by Mondrian about, or what was World War I all about, anyway? Even if we confine ourselves to linguistic utterances, we find ourselves with a slippery notion that is subject to all sorts of doubt and uncertainty. In saying to a person on the street "My garden is poor this year", I could very sensibly be talking about the cool weather, the lack of rain, the presence of pests, or a decision I made some time ago to plant a certain variety of tomatoes. I could be talking about any of these things, and more. However, the one thing that **is** clear that I am talking about in this instance that seems inescapable is quite simply that I am, in fact, talking about my garden. This is, obviously, because in uttering the sentence, I use the phrase *my garden*, whereas in this instance there is no particular mention of rain, weather, pests, or poor plant selection. To distinguish these two types of *aboutness*, the term **reference** is going to be used for those things overtly mentioned in the utterance of a sentence. Thus, I may be talking about the dry weather, but I am **referring to** my garden (and, the current year as well).

This is helpful, in that it localizes and objectifies a certain type of aboutness in a reasonably clear and intuitive way. Yet, even here there is

all manner of cause for question and uncertainty. For example, in the utterance above, might I also be referring to myself (by using *my*), gardens in general (by using *garden*), the quality of being poor, and so forth? Intuitively, these questions have sensible answers both yes and no. But it does remain a very solid intuition that I am referring to my garden, where an intuitively-based denial would seem far less convincing. For this reason, the focus of a theory of reference has been on those elements of a sentence or utterance which most clearly display the intuitive phenomenon of reference, leaving aside the subsequent questions for resolution within a more precisely articulated theory. The types of words and phrases that canonically display reference (see Strawson 1950) include demonstrative and indexical words and phrases (e.g. *this table*, *that cat*, *I*, *this*), proper names (*Aristotle*, *Paris*, *Fred Smith*), and singular definite terms (*the woman standing by the table*, *my garden*, *the author of* The Republic). Phrases and words of these types, not only in English but where they appear in any other natural language, unequivocally "pick out" some particular, definite individual or object. The point is, if these things don't exhibit the phenomenon of "reference," then we should all close up shop on this particular topic and find something else to work on.

2. Semantic Reference

2.1 Frege. Reference, then, is a kind of verbal "pointing to" or "picking out" of a certain object or individual that one wishes to say something about. But what, then, is the connection between the meanings of the particular words of the language we use in order to accomplish this, and what is picked out as a consequence? In order to frame this question, let us consider what has been typically called the NAIVE THEORY OF REFERENCE. This was by no means first articulated by Frege (1879) (one immediate precursor was Mill (1843)), but Frege seems to have taken the idea and pursued it the furthest within a new conception of how to do things--using the tools of formal logic--that appears to have been a genuinely novel development on the intellectual scene. The basic notion is that the meaning of an entire (declarative) sentence of a natural language is intimately connected to its truth value, and the contributions of the words and phrases within a sentence to the meaning of the whole are determined

by the contribution they make to the truth-value of the whole. Or, as McGinn (1981) puts it: "Reference is what relates words to the world of objects on whose condition truth hinges." (p157)

If one then turns specifically to intuitively referential phrases and words and calculate the contribution they make to the truth-value of the whole, one encounters and initially surprising result: that the truth-value of sentences containing referential phrases is (in part) determined by what the phrases themselves refer to, and not by any other or further characteristics of the phrases themselves. From an intuitive point of view, if I say (falsely) that Ringo Starr wrote the novel *War and Peace* then the truth-value of this sentence has not to do with any particular beliefs or conceptions I or anyone else might have about the world, but rather what Ringo himself, that guy out there, has and perhaps has not accomplished. Let ☺ be the person Ringo Starr. It is as if I am saying something to the effect that: ☺ wrote *War and Peace*.

Slightly more technically, and the success of this is easy to overlook, *any* phrase that has the reference ☺ will be automatically guaranteed to yield a sentence of the same truth-value if placed in the same syntactic location in the sentence as the phrase *Ringo Starr*. Thus, supposing that the phrases *the most famous drummer for the Beatles*, *Jimmy Smits' boyhood hero*, and *that man over there* have, on an occasion of use, the reference of Ringo Starr, their contribution to the meaning of any sentence will be ☺ and nothing more. This will mean that all the following sentences are likewise guaranteed to be of the same truth value as "Ringo Starr wrote *War and Peace*"; as will, in fact, any other way whatsoever of referring to the particular man ☺ , for this is the contribution any such phrase will make to the truth-value of the whole:

(1) a. The most famous drummer for the Beatles wrote *War and Peace*
b. Jimmy Smits' boyhood hero wrote *War and Peace*
c. That man over there wrote *War and Peace*

Thus we have an actual diagnostic for what is intended by the term *reference*, namely, preservation of truth-value. Consider the following to

see how this might go. I wonder if the word *someone* is a referring term, and on an occasion of use can be used to refer to Ringo. (This would seem intuitively plausible under certain circumstances. Suppose, for instance, I host a birthday party for Jimmy Smits and have invited Ringo as a surprise guest, and when Jimmy complains how the party is dragging I might say presciently, "Yes, but someone has yet to arrive!") Now consider the contribution to the truth-value of "someone" in the following:

(2) Someone wrote *War and Peace.*

The judgments here are not wholly secure, but most people who think about these things agree that what has been said here is, in fact, true, whereas if it were referential and had the value ☹, it would have to be false. Assuming these intuitions hold up, then *someone* is not a referential phrase (though see Fodor and Sag (1982), for a different point of view). It makes some other contribution to the meaning of the whole.

One might, thus far, look upon this discussion as a rearticulation of LEIBNIZ'S LAW of the intersubstitutibility of indiscernibles *salva vertitate*. But there are some objections to this that have been the source of continued inquiry to the present time, which Frege also tried to deal with, chiefly in Frege (1892). One objection, that I will mention and put to the side, is that one must not use examples where the use of a term is metalinguistic. Words and phrases function as names of themselves occasionally in language. When so construed, they do not have reference to the "usual" objects and individuals, but to **different** objects, e.g. the linguistic objects themselves. Thus (3a) is not to be intersubstituted for (3b) (preserving truth):

(3) a. Ringo Starr is a stage-name
b. The most famous drummer of the Beatles is a stage-name.
c. Richard Starkey's more famous alias is a stage-name

However, any other phrase with the reference the name *Ringo Starr* will preserve truth-value. Thus, (3c), unlike (3b), will have the same truth value as (3a), having the same reference.

While it is not always a straightforward matter to determine metalinguistic usage (e.g. consider the discussion of METALINGUISTIC NEGATION (Horn, 1989), this particular objection has had primarily nuisance value in the development of a theory of reference. More telling are one type of intuitive objection, and another based on failures of intersubstitutibility. The intuitive objection can be simply illustrated thus. If the meaning of a word or phrase is its reference, and "Ringo Starr" and "the Beatles' most famous drummer" have the same reference, then they have the same meaning. This just plain is not so: these phrases have obviously different meanings. This objection has clear force. The other objection gets to the heart of the naive theory of reference: that phrases with the same reference are not always intersubstitutible preserving truth-value. This phenomenon has received a huge amount of attention in the literature. One facet of this objection comes from the behavior of propositional attitudes. The following pairs of sentence can easily diverge in truth-value:

(4) a. James believes that Ringo Starr is a solo singer.
b. James believes that the Beatles' most famous drummer is a solo singer.

Having followed sporadically the later stages of Ringo's career, and having no idea whatsoever of any connection he might have had to the Beatles, James could well be described as having the first belief, but not the second (he assumes any such drummer is a drummer and not a solo singer). The problem is, if reference, under the naive theory, is all that contributes to truth-value, then both sentences could (crudely again) have the following contents:

(5) a. James believes that 😐 is a solo singer.
b. James believes that 😐 is a solo singer.

and so, being identical in contents, have the same truth-values. To object that there is, in fact, a reading of these sentences which does have this consequence--the *de re* reading where, intuitively, the speaker is the one

taking responsibility for the contents of the referring phrases--does not adequately address this point. There is a reading (perhaps the more natural one) where identity of truth-value is not the consequence, and on the purely naive theory of reference discussed here this simply should not happen. This is traditionally called the *de dicto* reading, and if the theory thus far is correct there should be no such phenomenon.

The other major type of consideration is that of the contents of identity sentences, which are generally assumed to be successfully analyzed by the "=" relation. Such sentences do not appear to introduce operators giving rise to opaque or *de dicto* contexts, but nevertheless are a similar source of puzzlement. If the contribution to the meaning is the reference of the noun phrases in the following sentences, then both ought to have the same "cognitive value" (a phrase that will be somewhat clarified below).

(6) a. The Beatles' most famous drummer is Ringo Starr.
b. Ringo Starr is Ringo Starr.

That is, both have the value:

(7) ☹ = ☹

But while the second is very obvious and can be known to be true a priori (assuming both instances of "Ringo Starr" are the same, see below for comments on this), the first seems to convey contingent information that may actually come as news to some people. This is a genuinely different kind of objection, because in fact '=' preserves truth-value given identical referents. Whichever way one finds of referring to Ringo Starr, intersubstitution will in fact yield identical truth-value.

Frege's proposed solution to these problems is well-known and often written about, but is itself problematic. The proposal is that words and phrases, besides having a reference, also have something which, in English, is called a SENSE. This "sense" of a word or phrase is what distinguishes otherwise coreferential expressions. *Ringo Starr* and *the Beatles' most famous drummer* may have identical referents, but are

distinguished by their senses. The sense contains the "mode of presentation" of a referent; it is an objective, and not a subjective thing, but it is what we psychologically "grasp" in understanding a word or phrase, and in so grasping enables us to find out the reference of the word or phrase. However, in Frege's view, it is not the psychological grasping itself that actually determines the reference, but rather the objective sense itself that is responsible for determining the reference. Thus, reference is determined indirectly from expressions of a language (this includes mathematical notation): a bit of language **expresses** a sense, which in turn **determines** a reference. This holds in the case of proper names as well—the names *Richard Starkey* and *Ringo Starr* (or *Hesperus* and *Phosphorus*, or *Cicero* and *Tully*, to revert to more traditional examples) have different senses associated with them despite common reference.

Frege's solution to the problem of *de dicto* meanings appears, initially at least, to work but strikes many people as unduly complex and counterintuitive (see especially Barwise and Perry (1983)). In certain syntactically-definable contexts such as embedded clauses, a referring expression does not have as its reference its "usual" one, but rather its sense. Thus, in the propositional attitude examples such as those above, the reference of *Ringo Starr* and *the Beatles' most famous drummer* is not the "usual" ☺, but rather the "customary" senses of each (we will call them S_1 and S_2), which differ from one another, and it is these senses which now contribute to the meaning of the whole. And since the senses are different, one now has different propositions that can diverge in truth value:

(8) a. James believes that S_1 is a solo singer.
b. James believes that S_2 is a solo singer.

There are two somewhat odd consequences of this solution. If the sense of an expression is, in these instances, its reference, and reference is determined by its sense, and since the reference determined by the senses S_1 and S_2 is ☺ in contexts such as those in (6), then there has to be *another* sense (S_3) that will determine S_1 and still another (S_4) that will

determine S_2 as their references in examples such as those in (8). But, unlike the customary senses, we have no clear intuitive grasp of what these might be. Further, the claim is that referring phrases in *de dicto* contexts have as their meanings different things from what they have in *de re* contexts. Complicating things still further is that if you have a *de dicto* context embedded within another *de dicto* context (e.g. "John was surprised that the Queen of England believed the Beatles' drummer was Ringo Starr"), then in the most deeply embedded context, the reference of a referring expression is no longer its customary sense, but rather the sense that determines its customary sense, introducing a third-order sense that must determine that as its reference. This works recursively, so that if there are n embedded contexts in a single sentence, in the nth context there would have to be an nth+1 sense to determine the nth sense as *its* reference. (This is not an incoherent proposal. Within the formal framework of Montague (1973), for example, what correspond to such higher-level senses are recursively definable, though any sense beyond the customary one is a constant function).

The oddness is compounded somewhat by Frege's view of sentence meanings. In ordinary contexts, such things do have a reference, which he takes to be a truth-value, and a sense, which he takes to be a proposition of a "thought". In *de dicto* contexts, however, the same recursive piling up of senses occurs as with referential phrases, so that an embedded sentence ends up *meaning* something different from its unembedded counterpart, a doubly embedded sentence has still another meaning, and so on. This strikes many people, again, as a bit strange.

When we return to the issue of identity sentences, which do not involve *de dicto* contexts, it is a little hard to see how Frege's suggestions lead to a definite solution. For the phrases used have their customary reference, so all true identity sentences express a proposition of the form **a**=**a** (where **a** is some arbitrary referent), though within a belief context, for example, different senses will emerge to distinguish the (higher-level) propositions created.

Making use of the discussions to be found in McDowell (1977) and Dummett (1975), this is what may have been intended. From the point of view of one understanding an utterance, "grasping" the sense, which determines the reference, does not enable one to automatically grasp the reference itself. If this were so, and we happened on Smith foully murdered, all we would need to do is to hear someone utter the phrase (in a *de re* context) *Smith's murderer* and the identity of the murderer would be automatically known to us; but, obviously, it is not. Likewise, we would (in at least one uncharitable interpretation of Frege's framework) only have to understand a sentence in order to know its truth value. To check on how many copies of an article we need to submit to a journal for publication, we'd only need to hear someone go through a list "*The Journal of Modern Fregean Studies* requires one copy...two copies..." etc. until we hit on the reference "true".

But there has to be some kind of psychological connection between grasping a sense and determining a reference (let's call this relation "finding" a reference, incorporating Russell's notion of "acquaintance" as a "direct cognitive relation" (Russell 1910). Consider, for instance, your understanding of the phrase "My sister's oldest daughter". If you can read this paper then you clearly understand what this means—you "grasp" its sense—but it is very doubtful you are antecedently familiar with that particular person. That is, the reference is unknown to you. However, if you wanted to go to the time and trouble to discover the identity of that person, the meaning of the phrase itself provides you with some kind of clue about how you could go about finding the reference. For instance, you could ask me who my sister is and how you can reach her, and then ask her or whoever answers the door who her oldest daughter is, and then go find her. On the other hand, if I used the phrase, which you understand, "The best young salesperson at the Anthropologie store located in downtown Seattle", to make reference to the same person, you'd likely skip hunting up me and my sister and head for the manager of that store in Seattle. In saying that understanding the contents provides one with "some kind of clue" about how to find the reference, I am not implying either that grasping the sense provides one with anything like a definite **procedure** for finding the reference, nor that it provides

one with any guaranteed **means** of making that identification. Furthermore, in this framework, the "clues" are not the meaning of the phrase (for the meaning is an objective, not a psychological matter), but rather psychological addenda that are intended to elucidate further the notion of "mode of presentation".

Let us return to the point at hand, identity sentences. While in the Fregean framework it appears one cannot make use of the notions of sense and reference directly to distinguish *de re* propositions with terms of identical reference, the "cognitive values" will differ. In asserting, for example, that Plato is the author of *The Republic*, different means of finding references are suggested by the phrases *Plato* and *the author of* The Republic, and the information conveyed by asserting such an identity sentence is that the clues provided by each via grasping their different senses will converge on the same reference. That is, there is differing **psychological** information associated with the use of each phrase, even if the contents of the proposition expressed is of the form **a**=**a**.

But, even if this view holds any validity, it generates a subsequent puzzle, for what is the value in Frege's framework of grasping the sense of a proper name? Beyond the fact that the reference is (at least occasionally) called by that name, and in the normal case many people are (at least occasionally) called by that same name who cannot be further distinguished, the sense of a proper name seems largely if not entirely devoid of any "clues" about how to determine their referents. Nevertheless, the very fact that the names differ suggests at least a partially distinguishing means of identification. We will see other suggestions about this later.

There have been other problems and questions raised by the Fregean framework that have been pursued in subsequent work. Modal contexts, for example, have been noted as providing similar puzzles to those posed by propositional attitudes. A particularly knotty problem for the framework is posed by negative existential sentences, such as "The king of France does not exist." Let ☺ be the king of France. If this is so, then there is, in fact, some king of France, namely ☺. Thus, a proposition of

the form ☺ *does not exist* would appear contradictory. Now suppose, as is currently the case, that there is no such ☺. The form of the proposition would appear to be an unsaturated proposition of the form __ *does not exist*, which is assigned no truth-value since it lacks anything that the phrase *the king of France* contributes to the proposition.

2.2 Russell. There were two attempts in the same era intended to resolve this particular difficulty. One was the "bite the bullet" analysis of Meinong (1904), who made what some find the curious claim that, in fact, phrases like *the golden mountain*, *the square circle*, and similar phrases, including proper names like *Zeus*, do in fact have reference. It's just that they have the property of not existing, but such phrases can and do contribute a reference to the proposition. Such a solution strikes many as ontologically a bit bizarre, for if there is no such thing as a golden mountain, then surely "the golden mountain" has no reference, there being no such thing. This seems transparent reasoning. At least that seems the attitude of Russell (1905), who proposed instead a different, and what people at least for some time considered a much more clever, and more ontologically satisfactory, type of solution—his theory of definite descriptions. This theory addressed the problems of *de dicto* contexts, and of identity, so there was a lot of mileage to be gotten here.

The strategy Russell followed was, in effect, to deny that definite descriptions (and proper names were taken as a variety of definite descriptions) contributed anything like a reference to a proposition. In fact, definite noun phrases did not contribute to a proposition any identifiable single constituent of meaning at all. Rather, these were disguised existential statements which were assertions of uniqueness. To illustrate, take the sentence "The Queen of England is dignified". This, setting aside analysis of the name *England*, cashes out as:

(9) ∃x [Queen of England (x) & ∀y [Queen of England (y) → y=x & dignified (x)]]

From there it is a short step to negative existential sentences, provided one has a syntactic means of according the negation widest scope. "The king of England does not exist" comes out as an unremarkable statement:

(10) $\neg\exists x$ [King of England (x) & $\forall y$ [King of England (y) $\rightarrow$ y=x]]

Proper names are taken as disguised definite descriptions and analyzed accordingly. Thus, *Pegasus* might have the contents of being a winged white horse of mythology, abbreviated as WWH, and "Pegasus does not exist" similarly comes out as:

(11) $\neg\exists x$ [WWH (x) & $\forall y$ [WWH(y) $\rightarrow$ y=x]]

There is no need for Meinongian non-existent objects, or any strange reference at all since, since definite descriptions and names do not have reference in the first place.

Identity statements like "Hesperus is Phosphorus" become, on this view, unproblematic as well. Let the contents of *Hesperus* be ES and that of *Phosphorus* MS. Abbreviating by omitting the uniqueness clauses for the sake of simplicity, the identity statement comes out something like a fairly ordinary looking assertion:

(12) $\exists x \exists y$ [MS(x) & ES(y) & x=y]

And, finally, the problem of *de dicto* contexts receives a treatment that avoids the piling up of senses that is a consequence of the Fregean analysis, for the propositions expressed have different forms (RS abbreviates whatever descriptive contents the name *Ringo Starr* has, and, again, the uniqueness clauses are omitted for simplicity).

(13) a. James believes that Ringo Starr is a solo singer.
b. James believes that the Beatles' most famous drummer is a solo singer.

(13') a. James believes [$\exists x$ [RS(x) & solo singer (x)]]

b. James believes [∃x [most-famous-drummer (x) & solo singer (x)]]

The *de re* reading would simply accord the existential expressions widest scope.

In Russell's framework, then, did anything at all have reference value? He did admit of something called a logically proper name which despite terminology is no proper name but would strike most as an indexical expression. This is exemplified by directly-referring demonstratives without nominal contents such as *this* and *that* used in a context to make unmediated reference to some individual or object. Note that such instances do not, in fact, cause immediate difficulties for identity, or for negative existentials (e.g. "This (said, pointing at a table) does not exist" seems a blatant contradiction), or for *de dicto* contexts.

This would appear to be a significant improvement, but this theory too has been met by influential reply on two major fronts. One, articulated by Strawson (primarily, Strawson (1950)) questions whether Russell's theory of descriptions might be missing something crucial. The other type of objection concerns the descriptive contents of proper names, discussed by Kripke (primarily, Kripke (1972)), which gave rise to the idea of **direct reference** theories of names. (Though for an updated defense of Russell's position, see Neale (1990) and Ludlow and Neale (1991)).

3. Reference as Pragmatic

3.1 Strawson. The fundamental question Strawson raised is whether what we are calling "reference" is a matter of (linguistic) meaning. Both Frege and Russell expounded what we can call a "semantic" theory of reference, in the sense that a semantics characterizes the meanings of words and phrases of a language in a general sense. Any further meaning that results from producing and understanding the actual utterance of a sentence, which are types of human actions, is not characterizable within the semantics as there is no reference there to speakers and hearers, only words, phrases, syntactic categories, etc. Thus, Strawson argues, truth and falsity are not (semantic) properties of sentences of a language, but

rather is a property of a **use** of a sentence (via an utterance) on a particular occasion. To illustrate his point, he presents the example of "I am hot". Now, he points out, it makes no sense to ask if this **sentence** is true or false; it is only when use is made of the sentence that we can so evaluate it. Thus, at this moment, if I pointlessly utter "I am hot" aloud, it's false, but said by another at this very moment it might be true. He points out, similarly, that a referring noun phrase like *the king of France* can only be evaluated for its reference value with respect to a use of the term via an utterance as well. It just so happened that Russell, writing in the early 20th century, was writing at a time when there was no king of France; had he written two centuries earlier, there would have been and the phrase would have had a reference. But again, this is not a fact about the noun phrase meaning itself, but about a particular use of the noun phrase.

He further points out that the verbs *mention* and *refer* (he treats them as synonymous) are verbs of **doing**. This point is elaborated on more clearly by Linsky (1963), who notes that the verb *refer* does not have only a general sense (e.g. as in "x refers to y"), but also a specific sense that may be applied to individuals, such as in saying, to use his example, "Who are you referring to when you say 'the Sultan of Swat'?", and receiving the reply "I am referring to Babe Ruth, of course". This more specific (non-stative, achievement verb) is not something we apply to language: it is strange to say "'The Sultan of Swat' is referring to Babe Ruth." Linsky points out that the question, to what does x refer? is a different question from asking, what are you referring to in using x? Further, in most instances, definite descriptions cannot be said, in general, to have any reference at all, even if they have meaning. Thus, "the man with the gray hair" is and can be used to refer to some particular man on a given occasion of use, and another on another occasion, but the semantics of this phrase does not pick out some unique individual (there being many with gray hair) *simpliciter*.

To revert momentarily to the Fregean framework, we might reason thus: If the contribution to the meaning of a proposition is the reference of a phrase, and we determine the contribution some expression makes to the

meaning of a proposition in terms of its contribution to the truth-value of the sentence expressing it, and if the notion of truth-value is something assigned to specific uses of a sentence via its utterance (and is therefore not a part of the semantics proper)—then, the notion of "reference" plays no (clear) role in the semantics proper but, rather, only in the use of a sentence on a given occasion.

Strawson points to another intuitive phenomenon that Russell's analysis provides no room for. If I were to say to you, right now, "The king of France is wise", on Russell's analysis I would have said something false (there being no such unique king). However, Strawson raised the point that it does not seem to be true or false *simpliciter*, but rather, there is something funny or strange about the utterance. His intuitions are, if asked if the sentence so used is true or false, he'd be at a loss. As he puts it, since there is no king of France, then the question of truth or falsity of the sentence simply does not arise. The use of the phrase *implies* that there is such a king, and in the absence of the validity of this implication one can make no sense of evaluating for truth or falsity. In more modern terms, this would be called PRESUPPOSITION failure. Note, again, that this "implication" can hold or not hold depending upon when the sentence is used, so its holding or not is thus a matter of usage, not of semantics proper. Strawson holds that even in cases of presupposition failure, the **sentence** still has "significance"—it is not gibberish—but this significance is not grounded in evaluating for truth or falsity. Strawson also goes on to point out that the same phenomenon occurs with Russell's logically proper names, which behave in this respect just like definite descriptions, despite having a very different analysis in Russell's theory.

3.2 Kripke and "direct reference". Kripke's critiques (Kripke, 1972) focus on something quite different—the analysis of proper names that Russell presented (and also, though perhaps somewhat unfairly, Frege presented in according names senses), in which names were treated as disguised definite descriptions. His critiques more or less interweave with similar critiques of Putnam (1975), Barcan Marcus (1963), Donnellan (1972), Geach (1962) and Kaplan (1986). Kripke argued very persuasively that any such analysis will fail, and that another

understanding of the nature of the reference of proper names is required in which names refer directly to their referents, without the mediation of any sense or descriptive contents (which is very similar to the "naive" theory of reference). The arguments take a variety of forms, so I will present only a couple here to illustrate how this conclusion might be reached. An excellent summary may be found in Salmon (1989).

Suppose that we provide the name *Ringo Starr* with the content "(most famous) drummer for the Beatles", or something similar. The modal argument is that this would entail that no matter what the circumstances, the drummer for the Beatles would be Ringo Starr. This is, if Ringo had not passed his audition, and someone else had, then that person would be Ringo Starr, not the unemployed drummer wandering the streets of Liverpool. This type of argument rests on contrafactual thinking, but seems fairly persuasive. That Ringo Starr is the drummer for the Beatles is not a logical truth. The epistemological argument has a similar flavor; it rests on the possibility of mistakes. This works best for historical figures around whom legends have developed, where the proposed meaning of a name may contain all sorts of factual error, and identify either no one at all, or by chance someone else, which seems counterintuitive.

The strongest argument though seems to be the semantic argument, which does not deal in possibilia but relies upon our judgments about who or what a name does in fact refer to. Donnellan (1972) uses the example of a person named Thales, referred to by Aristotle and Herodotus, among others, as a philosopher who held that all was water. Now, suppose that Thales' view was, say, misinterpreted by these other philosophers (it happens!), and he in fact held a much more subtle view which, in the end, could not be accurately so characterized. Imagine also that Thales had a not very good student who was a philosopher, about whom we know nothing whatsoever, who also misinterpreted his teacher and in fact did believe all was water. If the phrase "the ancient Greek philosopher who believed all was water" expresses the content of the name *Thales*, it does not pick out Thales but his student instead. Now, the question is whether the name *Thales* as used by Aristotle and Herodotus and passed along to

us in fact refers to that obscure student. Manifestly, it does not (examine the discussion immediately above for clear evidence). It refers to, well..., Thales. The view that results from this general line of thought is that proper names are expressions which refer directly to their referents, and there is no mediating sense or meaning which is employed to necessarily determine reference. Further, this reference is RIGID, in that a name picks out the same individual in all possible worlds. Putnam (1975) has employed arguments that are similar in thrust to argue that natural kind terms, such as *tiger* or *gold*, lack extension-determining semantic descriptive contents.

But this seems to leave us with the problems of identity sentences and the other issues Frege was struggling with. In partial answer, Kripke outlines an approach (which he himself does not characterize as a theory) that is critiqued subsequently by Evans (1973). This is the CAUSAL THEORY of names. The general idea is this: a speaker who uses a given name **A** will make successful reference to the individual it refers to ☹ just in case there is a reference-preserving chain of usage of **A** that extends back to ☹. Informally, at first there is some veridical naming or "dubbing" that initially fixes the reference of **A** as the individual ☹. When one of the dubbers uses **A** in the presence of a further person Dr. X, then Dr. X's use of **A** will pick out ☹ on the strength of the dubber's (secure) usage. This works transitively, so that if Dr. X talks to Prof. Y and uses the name, Prof. Y is thereby entitled to use A and will thereby refer to ☹ as well, and so on.

Each name, including different names for the same individual or object, will appear in different utterances at different times, and almost certainly involve many people who are familiar with one name but not the other. That is, the usage of the names will distinguish different casual chains. Thus, what appears to be different senses attached to different names, like *Mark Twain* and *Samuel Clemens*, are not senses but causally-determined chains of usage. And, roughly, in asserting that Mark Twain is Samuel Clemens the new information imparted is that the distinct causal chains associated with these two names will ultimately converge on one and the same person. The task of converting this line of thought into an

actual theory, however, is daunting; Evans (1973) argues that the causal theory is problematic and resorts to a notion of communal knowledge about the use of a name and the intention to use the name in accordance with that communal knowledge. In some respects, this is similar to the causal theory, in that it does not accord proper names a semantic meaning, but rather reference is achieved via the mechanism of social practice.

Thus far, we have seen a communal line of thought in which the notion of reference has, in fits and starts, become increasingly removed from being a purely semantic notion, and increasingly a function of human action and interaction. At this point, though, I wish to step back and ask the extent to which this particular direction is justified.

4. Semantic Reference and Pragmatic Reference

4.1. Some issues. One problem with talking about "reference" is that this word has an ordinary common meaning that we are trying to accord a consistently-used technical meaning. Not only must we contend with this, but we also must contend with the extent to which translation decisions regarding Frege's use of *Sinn* and *Bedeutung* are best thought of as "Sense" and "Reference", when the German terms have no exact counterparts in English. Would things be different were these labeled "Concept" and "Designatum" instead? Perhaps—just perhaps—Frege was not even trying to talk about something called "reference" at all. We certainly want our discussions of reference to be substantive and not terminological. Thus far, we have been following the intuitive idea that the "reference" of a word or phrase is what contributes to the truth or falsity of a sentence (or a sentence in use). The strategy of fixing on those elements of language, such as names and definite descriptions, which most clearly are in line with our native understanding of the word, is an understandable and productive strategy. But what about all those other words and expressions in a sentence? Does a notion of "reference" fit them as well?

If we apply the notion, in slogan form: "reference determines truth", across the board, and consider all those other words and phrases in a sentence, it turns out that all of them play a role in determining truth and falsity. This, in very rough form, is what Church (1943) and Gödel (1944) quite independently noted about Frege's theory (see also Salmon (1981, 1986)). The possibility of syncategorematic introduction of some words to the side, what this means is that **all** words and **all** phrases within a sentence have a "reference" besides a "sense". If we examine, say, Montague's (1973) framework, which to a large extent can be thought of as one implementation of Frege's framework, at least in many important respects, most of the references of words and phrases turn out to be functions. This holds true as well for proper names (which turn out to be generalized quantifiers, that is, functions from properties to truth-values) as well as definite descriptions (to which he gives a type of Russellian analysis). What then becomes of our native notions of "reference" that lead us to see the sense in focusing on names and definite noun phrases to the exclusion of everything else? Again, our native notion of reference brings us to the underlying intuition that reference is a "picking out" of a definite object, which we want to say something about. And the type of linguistic device that seems best suited for these purposes are (certain types of) noun phrases. An understanding of why this might be so might be gained from an analysis of the variety of noun phrase meanings.

Let us examine a proper name, within the Montague framework. Although its "reference" within a sentence is the set of properties that individual might have (an analysis suggested in Leibniz (Mates 1968)), (nearly) any proposition in which this occurs is equivalent to another formally distinct proposition in which the individual itself appears as the subject (object, etc.) of predication. Thus, although the analysis of something like "John snores" is:

(14) $\lambda P[\check{}P(j)](\hat{}snore')$

Which is semantically equivalent to:

(15) snore' (j)

However, with quantificational noun phrases, which are not considered referential on the whole, no such similar reduction is possible. Since Montague chooses to represent definite descriptions in the Russellian manner, they too come out as quantified noun phrases. But updating the framework some, and allowing for presupposition failure with definite noun phrases, one can arrive at an interpretation which is of the form P(a), with a, as Russell might say, as the "logical subject" of the predication. In this way, we might begin to understand why certain types of noun phrases can be used "referentially", and others not. However, looking at things this way may seem to open the door once again to a semantic view of reference.

One view is discussed in McGinn (1981) and rests on an analogy. We do, in fact, make reference to things by uttering certain noises under certain circumstances. However, consider the commonplace activity of buying something. This might at first sight seem to be a causal interaction describable in such terms, but this doesn't seem correct (e.g. that we can buy things, and often do, which do not exist at the moment of purchase). In making a purchase, we operate against a background of conventions and constructs, such as the notion of money, ownership, legalities of exchange, and so forth. These notions are not behaviors, but rather collectively define certain types of economic relations between individuals and objects. In actually making a purchase, we are in fact guided by perception and perform actions, but these things are not part of the economic relations themselves. And we also have considerable freedom of behavior in that there are many ways to make a purchase. Rather, making a purchase results in a relation between us and objects that we "get into" by doing certain things in concert with others. Similarly, McGinn invites us to consider, reference is like this. Reference is defined semantically by a relation between expressions and objects of the world, but this relation is not defined in any way by actions or intentions. Rather, it is a relation we can "get into" with an object by acting certain ways in a given context.

There is also a concrete strategy for "semanticizing" reference. It is, basically, Kaplan's notion of how to accord a meaning to indexical expressions (see also the contribution on indexicals, this volume). Before doing this, though, we need to give a bit of background on the notion of "truth".

The focus on truth and falsity as indicators of the reference of the use of a phrase needs to be distinguished from the notion of giving truth **conditions**. Let's approach it this way, within a possible worlds framework. In saying something like "Larry is in Spain," I am expressing a proposition p. That proposition p is a function in some theories from possible worlds (and times, but I'll omit this) to truth-values. That is, given a certain possible world w, then p(w) will yield either T or F. At the moment, I happen to know that who I am referring to by using the name *Larry* is, in fact, not in Spain. That is, I know that p(w) = F. Does this mean that, as a speaker, I know which possible world w is? No, I do not. For instance, I do not know if this is among the possible worlds in which Larry is in Spain **and** the rug in my office was installed in a month ending with *-ber*. I have no idea when it was installed. What this means is this: I do not know which possible world is the actual world. Nor does anyone else: no one is in command of all the statements that are true and false in our world. This does not stop me from referring to it, indexically or by description; it's just that the world we happen to inhabit as the actual one is not identifiable to us as w^{338} rather than, say, w^{784}. We believe, of course, that we can narrow down the host of candidates—this is not one of the worlds where "Cleveland is the capital of France" is true, and a lot of other such things. This would seem something of a step in the right direction, but this assumes none of us hold any false beliefs about the world, nor that we hold any contradictory beliefs about the world, which only makes things worse for any kind of psychological identification. In short, if we attempt a psychologization of the possible worlds framework, our limitations prevent us from ever "homing in" on which world among them all we inhabit. This is not a quirky feature of a possible worlds point of view: any theory of intensionality leads to the same conclusion.

None of this means we can have no cognitive "grasp" of truth-conditions, however. For any world w^n, $p(w^n)$ is true or false. We just don't happen to know if it's our world or someone else's.

4.2 Kaplan's analysis. Let us, in this setting, move on to Kaplan's notion of "character" of indexicals. There is a persistent intuition that although a noun phrase like *this woman* can be used on various occasions to refer to different women, and the class of women so referred to has nothing qualitatively in common (apart from being referred to in that way), nevertheless the phrase *this woman* has a constant meaning that transcends its particular uses.

Kaplan (1986, 1989) proposed that indexical expressions such as *I* or *that man* are expressions which, due to their nature, are assigned extensions (or references) only when their interpretations take arguments that most other words (such as *man*) need not (though see below for some qualification). These arguments, or parameters of interpretation, are the CONTEXT: who is speaking, who is the addressee, where the speaker is located, who or what else is in the immediate environment, and so forth (see also Lewis 1979, Stalnaker 1978, and King 2001). The CHARACTER of an indexical expression, then, is that function from contexts to intensions; that is, from (say) world-time pairs to extensions (i.e. references, in the generalized Fregean sense of the term). What this buys us is a single meaning for indexical expressions, that is, as expressions which are characterizable by a single function.

What role, then, does context play in the case of expressions like *eat* or *man*? One can say that they are "sensitive" to contextual parameters, in which case they represent constant functions--that is, for any set of contextual parameters given a certain world-time pair, the result will be the same in all instances (Chierchia and McConnell-Ginet 2000). Or, one could propose that such expressions simply do not take such arguments. When we turn to indexical expressions, though, things become slightly more uncertain. Are they "insensitive" to world-time parameters of interpretation? It would seem in general that this is so--they have a rigid

designation nature to them. But this is not always the case, it would appear.

(16) a. Always do **today** what you could put off till **tomorrow**.
b. (Said, frustrated at where I put my glasses, again) Darn, they always seem to be over **there**, never over **here**!

Examples such as this seem to have natural interpretations that are at least time-dependent: always do things the day you are experiencing, and not the next; my glasses are always where I am not. If one takes these as serious data (some do not), then world-time pairs would appear to play a role in the case of at least some indexical expressions.

Is there a rationale for proposing that meanings have this two-stage nature to them? Is there any rationale for distinguishing between contextual parameters, and possible worlds (and times) in the construction of meaning? There may well be: one can be directly acquainted with the contextual parameters in a way that one cannot with worlds (and, I'd actually argue times as well, but have not). Context, unlike a possible world, is cognitively accessible.

4.3 Pointing? There is another question about whether the nature of reference is found in human action. In Grice (1975) the notion of CONVERSATIONAL IMPLICATURE is introduced. There, he takes pains to point out that such implicatures are a special case of much more general principles appearing to govern human social interaction in general, verbal or non-verbal. A similar point is made in the case of speech acts (Austin 1962, Searle 1969): that the act of congratulating someone verbally, for example, can be alternatively accomplished non-verbally; one can equally well threaten verbally and nonverbally. From this perspective, we might ask the question that, if reference is something we do, then how might we accomplish the same thing non-verbally? There is a fairly obvious potential answer, namely, the phenomenon of pointing.[1] But, does pointing constitute reference? Let us use the symbol ☞ to stand for the act of pointing at an object. It is quite clear that this gesture itself cannot function as a part of the meaning of a sentence the way a linguistic

expression can, for something like "☞ is green" at best, has a "word-play" nature to it, much like threatening to hit someone over the head with a newspaper and saying "I'm going to... (action of raising a rolled-up newspaper as if to strike the addressee in the head)", or the bumper stickers reading, "I my dog".

Another way of putting it is the question of whether pointing has semantic significance. Reimer (1991) argues contra Kaplan (1989) that this is so, at least in cases where pointing is an accompaniment to the utterance of a demonstrative indexical; it is not just an attention-directing device. Reference is, Reimer notes, coupled with a certain intention on the part of the speaker to make reference. But the intention does not always determine what is said. For example, if there is a single dog sleeping among felines and there is no accompanying demonstration, the descriptive contents of the phrase *that dog* will be sufficient in the context to uniquely pick out (or discriminate) a unique reference, regardless of speaker's intentions. Likewise, if there are multiple dogs but one is especially salient in the context, say one wildly barking dog among other sleeping dogs, reference is again secured by uttering *that dog* in context. And again, the claim is, that reference would succeed independently of speaker's intentions; a pointing demonstration would be redundant. But suppose there are two equally salient dogs in a context, say both barking wildly. Here, pointing does have semantic significance in that it discriminates one dog from another. The dog pointed at is thereby the reference, and whatever is said of it is, truly or falsely, said of that individual, again regardless of speaker's intentions. If the speaker intends to point to refer to the white dog, but instead points at the black one, the black dog and not the white one determines the truth or falsity of what is said.

4.4 Intentions. Bach (1992) argues, on the other hand, that the "best of intentions" are good enough; that is, demonstrations do not have semantic significance. He considers a scenario where there are two (equally salient) sets of keys on a desk, and the speaker says, "these keys are mine," but by simple error grabs her office-mate's keys. Here, the mistaken demonstration (or, rather, what directs the listener's attention),

it is argued, does not affect reference. The intentions of the speaker to refer to her own keys still holds. However, referential intentions are not just any intentions, but rather ones whose distinctive feature is that its fulfillment consists in its recognition. It involves also intending that one's audience identify something as the reference. This is the piece that is missing in the office keys example.

One's reference, in this case, is not fixed by one's beliefs; it is fixed by the intention to refer and the intention that it be recognized as such. Consider the case, Bach suggests, where the speaker is sitting at her desk where behind her on the wall is normally placed a picture of Carnap. Unbeknownst to the speaker, someone has replaced it in the night with a picture of former U.S. vice-president Spiro Agnew. In gesturing towards the picture and saying, "that is a picture of the greatest philosopher of the 20th century", the speaker's belief is that the picture is one of Carnap, but the intention is to make reference to the picture (which is of Agnew) on the wall and the hearer is to recognize this intention via the gesture. The speaker has said something that is false, despite beliefs that she might have been saying something true (or, at least less manifestly false). The intention to refer to a picture of Carnap is not the relevant intention here. It is instead the one that is made available and intended to be made available to the listener, which leads to the picture of Agnew which is in fact there. Since having intentions and intending are not a part of language itself, but a property of speakers and listeners that may be carried out or indicated by action, pointing has no semantic significance but functions as a highly reliable indicator of the right type of speaker intentions, which alone secure the reference. This is the thrust of Bach's argument.

This idea is at least partially supported by Donnellan (1966) in his article on definite descriptions. First of all, he distinguished between different meanings or uses (there is equivocation) of definite descriptions, the REFERENTIAL and ATTRIBUTIVE instances. The attributive uses seem to function in a way reminiscent of Russell's theory of descriptions. That is, the contents of the noun phrase determine the propositional content. This is clearest when the identity of the one referred to is unknown to

both speaker and hearer, as in happening on a colleague murdered and saying something about Smith's murderer, whoever that may be. There is, however, the referential use which does not rely heavily, (or maybe even at all?), on the contents of the noun phrase. This is most clearly the case when there is a mistake in the description. Donnellan asks us to imagine a circumstance where there is a man in a group of people at a party who has a drink in his hand, which is in fact water. But it looks like a martini to me, from this distance. If I say to someone: "The man holding the martini is a famous author", Donnellan argues, my reference to that particular man will be successful, and the truth or falsity of what I have just said will depend on whether that man, who is not holding a martini but a glass of water instead, is a famous author or not. My reference to that man was successful because the listener understood my intentions to refer to that person, whether cognizant of my misattribution or not. Further, suppose there is another man right next to this person who is holding what looks to me to be a malted milkshake, but suppose it's in fact a martini. Although my description would fit this other man, and attributively pick this other person out, the "referential" use that I intended does not, and it is my successfully executed intentions, recognized by the hearer, that determines propositional contents.

Birner (1989) questions whether the conclusion is appropriate. Imagine that Mr. Smith is brutally beaten, seriously but not life-threateningly. As he lies in an alley, along comes a kindly doctor who believes (falsely) that he cannot survive, takes pity on his suffering, and administers a lethal injection to end his pain. The doctor leaves, believing himself to be a good samaritan, when in fact he has just committed murder. Along comes our person-in-the-street who utters the line: "Smith's murderer is insane." Now, to whom is the speaker referring?

Under Donnellan's account, this can't be a referential use because the speaker has no idea who did the beating. He means "whoever it was" who did the beating—the classic diagnostic for the attributive usage. But according to Donnellan, the attributive use would "pick out" the doctor as the referent, because the doctor is the one who satisfiesthe description.

so there's no way a speaker can refer to whom he intends to refer to (i.e. the person who administered the beating) under Donnellan's account.

4.5 Meaning imparted via usage. Recall that on the direct reference theory of names, there is no propositional contribution of the interpretation of the name corresponding to a sense—there is only the reference. This left hanging, as noted above, the question of why we find "Sam Clements is Mark Twain" different in contents from "Mark Twain is Mark Twain". The causal theory was one attempt to begin to characterize this difference. Salmon (1991) takes this issue up in detail (along with propositional attitude contexts), and articulates a very interesting version of the "bite the bullet" approach. His view, stated broadly, is that using certain words or phrases not only gives us information about the contents of the expressions themselves--a proposition--but also about the intentions and beliefs of the speaker. There is a distinction between the information **contents** of a sentence, and the **information imparted** by its use in context. It is often observed that an utterance in context conveys information about a variety of matters that are characteristics of the act itself, and not part of literally what is said. Such things as distance from the listener, gender, emotional state, regional or foreign accents, etc. are a part of this information conveyed, as is illocutionary information: "I'll be there tomorrow" could be a prediction, a threat, a promise, an offer, depending on context. Some of this information is very hard to intuitively distinguish from what one might call core semantic contents. (The literature on SCALAR IMPLICATURES (Fauconnier, 1975) I believe illustrates this point quite well.)

What Salmon argues is that, yes, the propositional contents of names is just the object referred to itself, as direct-reference in its starkest form would demand. Thus, "Hesperus is Phosphorus" is in fact a proposition of the form **a=a**. Then why does it seem informative? This goes back to a fundamental criticism of Frege's notion of sense (*Sinn*). Frege thought (for the most part) that the sense would be that which determines the cognitive significance of an expression. He made no distinction, however, between the conceptual contents of an expression and how the reference (*Bedeutung*) of an expression is **secured**. This conflation, it is plausible

to think, led to his attributing a sense to proper names (which was a change from an earlier more direct-reference view he held). However, once we make this separation, it becomes clear that "securing" the reference is a matter of how actions are carried out and the context in which they occur, and in particular by using certain words. In using a name, we also impart the information that use of that particular linguistic form can, in context, secure a particular reference. If we use a different name, even for the same object, then we are imparting a related but *different* type of information. Thus a sentence like "Samuel Clemens is Mark Twain" will impart information that "Mark Twain is Mark Twain" will not, even if the contents are identical. Soames (2001) mounts a substantial defense of (approximately) this point of view.

Perhaps, in the end, a notion of "reference" as a type of direct connection to objects in the world might well be appropriate for both a semantics and a pragmatics, as Kripke (1977) suggests, and that terminology could be modified to distinguish them.

5. What can we refer to?

Work on reference has, as noted above, tended to focus on what one might call the clearest cases. But the boundaries have been, for the most part, fairly limited. There remain a great number of questions and issues that the direct-reference theory raises. One particularly difficult issue that has received a great deal of attention, though, is the simple question: if a referring expression has no object for its reference, then can it have any content? That there should be some contents is intuitively clear from consideration of names of fictional characters. We need to distinguish cases of fictional reference from failure of reference. After all, Superman does, in fact, wear a cape, and the Lone Ranger a white hat, and decidedly not the other way around; however, at this moment, if I use the term *the giraffe in my office* and there is no such thing, this is simply reference failure. As there is no Superman or Lone Ranger, either, how can we attribute truth-values to such sentences?[2] A number of answers have been suggested, such as Bertolet's (1984) notion that the content of such assertions is not the apparent subject, but that the myth or story (which is

a part of the real world) exists and is structured in a certain way. In saying, for instance, that Pegasus does not exist, one presupposes that there is a story in which Pegasus figures, but asserts there is no real individual which the story is about. Hintikka (1983), on the other hand, takes a possible worlds approach and suggests that Pegasus, Superman, and the others are to be found as objects, but in other possible worlds from our own. So when we make reference to them, we are doing so in those worlds where they do exist, just not this one.

But "nonexistent objects" are not limited by any means to fictional characters (Parsons 1980). There is the case of dreams and hallucinations, for instance, where we use referential terms to make something like successful reference. If, for instance, while in therapy I hallucinate that there is a wolf in the corner, the psychiatrist can ask me questions about it, such as what color it is and how big it is, and I can answer using the expression *the wolf*, even under circumstances where I know that I am hallucinating. We also, curiously, make reference to objects and events of the real world which very clearly do not exist. If I say "the earthquake is supposed to be here in three hours", there is no earthquake I am directly referring to, and there may never be one. Or, if we talk about the government's projected surplus, or the party planned for next Saturday, and none eventuate, what have we referred to? If I buy a bookshelf and a shelf is missing, I refer to the missing bookshelf but have no guarantee that one is, in fact, out there laying around somewhere. Maybe one was never manufactured in the first place.

Beyond this, there is the question of what types of things there are around that can be referred to. Let us ask, for example, about the limits of proper names. Quite obviously, we have names for people and pets, places (Antarctica) and buildings (Dewey Hall) and books and many works of art (*Pride and Prejudice*; *Venus de Milo*). Species names such as *canis domesticus* also appear to exist. And if this is a name, what does it name that is a particular in the world (a species?), and then why wouldn't *the dog* (in the sense of "The dog is a mammal") also be a species name? Putnam and others have suggested that natural kind terms, like *gold*, directly refer to natural kinds; if so, would *honesty* be the name of a

characteristic or property? Is *spring* the name of a season (after all, *August* would seem the name of a month)? The question here is, if a hallmark of proper names is that they name particulars that they take as their referents, then it is possible we have a lot more particulars around than we might have thought, or that things other than particulars can be made direct reference to. The same might be said of definite descriptions. If I talk about "the average French voter", which is no particular French person, am I making something we can meaningfully call *reference*? If, as Chomsky notes, there is a flaw in my argument, and I refer to the flaw later on (in correcting it, for instance), does this mean there are such things as flaws out there as particulars?

Finally, and this could be related to the question of fictional names, the use of referential-sounding expressions can be used for entities only discussed with respect to a particular stretch of discourse (Kamp and Reyle (1993) call this "spontaneous fiction"). Many instances of this can be found in the above discussion: Suppose there is a man at the door. ***That man*** knocks. I talk to **him**. **He** introduces **himself** as Steve. I ask **Steve** to leave. **The guy** was at the wrong door.

What am I referring to in these instances? There is no such person, yet I have used what appears to be ordinary referential language to talk about him. The solution in the framework of Kamp and Reyle (1993) is to suggest that whenever we have (apparent) reference, we are in fact positing discourse markers and associating predicates with them. That is, the reference of language is not objects in the real world, but rather DISCOURSE ENTITIES. It is another stage of evaluation at which there is anything like a semantic mapping to the world (in this instance, models), where distinct discourse entities, for example, may or may not be mapped onto a single real individual. This is hardly a direct-reference way of looking at things, but holds out interesting possibilities.

Greg Carlson
University of Rochester

Endnotes

1. By "pointing", I intend any means of indication of an object, whether one uses the index finger, an open hand, a sideways nod of the head, one's lips, chin, etc. What constitutes "pointing" is, like other gestures, subject to cultural variation.

2. It seems likely that no language distinguishes the fictional from the real in terms of reference. Gregory Ward notes that he made an inquiry on the Linguist List as to whether any known language formally distinguishes noun phrases making reference to fictional entities and real ones, and none were reported.

References

Austin, John L. (1962). *How to do things with words.* Oxford: Clarendon Press.

Bach, Kent (1992) Paving the road to reference. *Philosophical Studies* 67, 295-300.

Barcan Marcus, Ruth (1963). Modalities and intensional languages. In *Boston Studies in the Philosophy of Science, Vol. I.* Dordrecht: Reidel. 77-116.

Barwise, Jon, and John Perry (1983). *Situations and attitudes.* Cambridge, MA: MIT press.

Bertolet, Rod. (1984). Reference, fiction, and fictions. *Synthese* 60, 413-37.

Birner, Betty (1989). On the 'referential-attributive' distinction, in Birner, B. et al (eds.), Working Papers in Linguistics, Northwestern University. 2:1-12.

Chierchia, Gennaro and Sally McConnell-Ginet . (2000). *Meaning and grammar* (2nd edition). MIT Press.

Chomsky, Noam (1981) *Lectures on Government and Binding* Dordrecht:

Foris.

Chomsky, Noam (1992) “Language and Interpretation: Philosophical Reflections and Empirical Inquiry,” in J. Earman (ed.) *Inference, Explanation and other Philosophical Frustrations* (Berkeley: Univ. California Press). 99-128

Chomsky, Noam (1995) “Language and Nature.” *Mind* 104: 1-61.

Church, Alonzo (1943) Review of Carnap's Introduction to semantics. *Philosophical Review* 52, pp. 298-304.

Donnellan, Keith (1966) Reference and definite descriptions. *Philosophical Review* 75: 281-304.

Donnellan, Keith (1972). Proper names and identifying descriptions. In D. Davidson and G. Harman (eds.), *Semantics of natural language.* Dodrecht: Reidel. 356-379.

Dummett, Michael 1975. What is a theory of meaning? In Samuel Guttenplan (ed) *Mind and Language*. Oxford: Clarendon Press. 97-138.

Evans, Gareth (1973). The causal theory of names. *Aristotelean Society*, supl. vol. 47. 187-208.

Fauconnier, Gilles (1975). Pragmatic scales and logical structures. *Linguistic Inquiry* 6, 353-375.

Fodor, Janet & Ivan Sag (1982) “Referential and Quantificational Indefinites” *Linguistics and Philosophy* 5: 355-398

Frege, Gottlob (1879). Begriffschrift. Reprinted in M. Beaney, (ed), *The Frege reader* (1997). Oxford: Blackwell. 47-78.

Frege, Gottlob 1892. Über Sinn und Bedeutung. Zeitschrift für Philosophie und Philosophische Kritik 100. Reprinted as "Sense and Reference" in P. Geach and M. Black (eds.) (1952) *Translations from the writings of Gottlob Frege*, pp 56-78. Oxford: Blackwell.

Geach, Peter (1962). *Reference and generality.* Ithaca, NY: Cornell University Press.

Gödel, Kurt (1944). Russell's mathematical logic. In P. Schilip (ed), *The philosophy of Bertrand Russell.* New York, Tudor Publishing.

Grice, Paul (1975). Logic and conversation. In P. Cole and J. Morgan (eds) *Syntax and semantics*, Vol 3. New York: Academic Press. 41-58.

Hintikka, Jaako (1962). *Knowledge and belief.* Ithaca, NY: Cornell University Press.

Hintikka, Jaako (1983). Situations, possible worlds, and attitudes. *Synthese* 54, 153-162.

Hockett, Charles, and S. A. Altmann (1968). A note on design features. In T. Sebeok (ed.), *Animal communication.* Bloomington: Indiana University Press.

Horn, Laurence (1989). *A natural history of negation.* Chicago: University of Chicago Press.

Hornstein, Norbert (1984) *Logic as grammar*. Cambridge: MIT Press)

Kamp and Reyle (1993). *From discourse to logic*. Dordrecht: Kluwer.

Kaplan, David (1989). Demonstratives. In J. Almog, J. Perry, and H. Wettstein (eds), *Themes from Kaplan*. 481-563. Oxford: Oxford University Press.

Kripke, Saul (1972) *Naming and necessity*. Cambridge: Harvard University Press.

Kripke, Saul (1977). Speaker's reference and semantic reference. . In P. French et al (eds), *Contemporary perspectives in the philosophy of language,*. Minneapolis: University of Minnesota Press. 6-27.

Lewis, David (1979) Scorekeeping in a language game. *Journal of Philosophical Logic* 8, 339-59.

Linsky, Leonard (1963). Reference and referents. In C. Caton (ed). *Philosophy and ordinary language*. Champagne: University of Illinois Press. 74-89

Ludlow, Peter and Stephen Neale (1991). Definite descriptions: in defense of Russell. *Linguistics and Philosophy* 14, 171-202.

Ludlow, Peter. (2003) "Referential Semantics for I-Languages?" in N. Hornstein & L. Antony (eds.) *Chomsky and his Critics*. Oxford: Blackwell.

Mates, Benson (1968). Leibniz on possible worlds. In B. van Rootselaar and J. F. Staal (eds.), *Logic, methodology and the philosophy of science III.* Amsterdam: North-Holland.

McDowell, John (1977) On the sense and reference of a proper name. *Mind* 86 159-85.

McGinn, Colin (1981). The mechanism of reference. *Synthese* 49, 157-186.

Meinong, Alexius (1904). *Untersuchungen zur Gegenstandstheorie und Psychologie*. Leipzig: Johann Ambrosius Barth.

Mill, John Stuart (1843). Of names. In: *A system of logic*. Book I, ch. 2.

Montague, Richard (1973). On the proper treatment of quantification in ordinary English. In J. Hintikka et al (eds.), *Approaches to natural language*. Dordrecht: Reidel. 221-242.

Neale, Stephen (1990) *Descriptions*. Cambridge, MA: MIT Press.

Parsons, Terence (1980). *Nonexistent objects*. New Haven: Yale University Press.

Putnam, Hilary (1975). The meaning of 'meaning'. In Putnam, H. Mind, language, and reality: Philosophical papers vol. 2. Cambridge: Cambridge University Press. 215-271.

Reimer, Marga (1991). Do demonstrations have semantic significance? *Analysis* 51, 177-83.

Russell, Bertrand (1905). On denoting. *Mind* 14, 479-93.

Russell, Bertrand (1910-11) Knowledge by Acquaintance and knowledge by description. *Proceedings of the Aristotelean Society* 8. 108-128.

Salmon, Nathan (1981). *Reference and essence.* Oxford: Blackwell.

Salmon, Nathan (1986). *Frege's puzzle*. Cambridge, MA: MIT press.

Salmon, Nathan (1989). Reference and information content: names and descriptions. In D. Gabbay and F. Guenthner (eds) *Handbook of philosophical logic.* Dordrecht: D. Reidel. 409-461.

Searle, John R. (1969). *Speech acts: an essay in the philosophy of language.* Cambridge: Cambridge University Press.

Soames, Scott (2001). *Beyond rigidity: the unfinished semantic agenda of* Naming and Necessity. Oxford: Oxford University Press.

Stalnaker, Robert (1984). Assertion. In P. Cole (ed) *Syntax and Semantics* 9. New York: Academic Press. 315-332.

Strawson, Peter (1950). On referring. *Mind* LIX, 320-44.

Wittgenstein, Ludwig (1953). *Philosophical investigations.* New York: Macmillan.