6
Referential Semantics for I-languages?[1]

PETER LUDLOW

Introduction

It is widely assumed that Noam Chomsky is hostile towards semantics, and at least in the oral tradition, it is supposed that this hostility carries over to virtually everything that goes by the name "semantics," including lexical semantics and model-theoretic semantics. This latter supposition is almost certainly mistaken, and it appears to reflect a confusion about precisely what kinds of semantic theories Chomsky objects to. For reasons that we will look at later, Chomsky has no problem with lexical semantics or model-theoretic semantics (at least in principle). Even when we turn to the domain of "referential" semantics – broadly construed as theories which articulate language/world connections – there are some versions that Chomsky will find unobjectionable.

My task in this paper will be to get as clear as possible on where Chomsky takes exception (and why he does). Surprisingly, we will find that once the landscape is properly surveyed, Chomsky's views do not exclude the kind of referential semantics that many philosophers favor. Indeed, I think that where there is disagreement, it stems not from disputes about referential semantics itself, but rather from certain subsidiary assumptions which I will try to make explicit and, in certain cases, defend.

1 Some Important Distinctions

To get some traction here we need to lay out some of the notions we will be working with. First, we need to get clear on the kinds of semantic theories that are in play (including those that might be characterized as being "referential" in some sense or other). Then we will need to get a bit more clear on Chomsky's notion of I-language.

1.1 Three Kinds of Semantic Theories

The first thing we need to understand is that "reference" is not a pre-theoretical notion, but as Chomsky (1995) has stressed, is a theoretical term the meaning of which will depend upon the broader semantical theory. As a very crude taxonomy, we can distinguish three ways in which "reference" is employed. As we will see, this three-way distinction still underdetermines the range of positions in play, but it will help us to clarify the subsequent discussion. I'll distinguish the senses as R^0, R^1, and R^2.

Many philosophers would take R^0 to be a misuse of the term "reference." According to this notion, reference does not involve relations to the external world, but is primarily a relation to internal representations. This might involve "reference" to elements of the "domain D" of Chomsky (1981), or to some sort of data structure or representation in the mind.

The R^0 notion of reference is fairly commonplace in linguistic theory (e.g., in lexical semantics like those in Levin and Pinker 1992) as well as in artificial intelligence research. Philosophical theories which employ the R^0 notion of reference include the structural semantics of Katz and Fodor 1963, Katz and Postal 1964, and more recent work in this vein such as Jackendoff 1972, 1983.

The key idea behind all of these proposals has been the notion that providing a semantics for a natural language expression (or syntactic form) requires that one provide a mapping of that expression (form) onto some representation which in some sense encodes the meaning of the expression. So, for example, in Jackendoff 1972, Katz and Fodor 1963, and Katz and Postal 1964, the idea is to map a given natural language expression onto a "semantic marker," which is in turn a symbol in a particular representational language (following Lewis 1972, we might call this language "Semantic Markerese").[2] In turn, Semantic Markerese is designed so that ambiguous expressions of a given natural language are disambiguated in Markerese. In addition, entailment relations and synonymy relations between natural language expressions are supposed to follow by virtue of the forms of their Semantic Markerese counterparts.

I would argue that the R^0 notion of reference is also employed by the various versions of model-theoretic semantics grounded in the work of Montague (1974). The idea that model-theoretic semantics cannot deliver language/world connections is controversial, and a number of model-theoretic semanticists are under the impression that language/world connections are precisely what they are offering (see Dowty, Wall, and Peters 1981; Partee 1980), but as Higginbotham (1990) and Lepore (1983) have argued, this betrays a certain confusion about the nature of model-theoretic semantics. Their argument is that model-theoretic semantics fails in exactly the same way that structural semantics is supposed to fail. What we expect a semantic theory to deliver is a characterization of the speaker's knowledge about the connection between language and the world.

Structural semantics does not deliver this because it only gives us a mapping from one language into some other system of representation. Model-theoretic semantics likewise fails, because it never makes the connection between language and the world, but at best gives us a mapping from expressions of a language onto certain model-theoretic objects.[3] This paper is not the place to take up the debate about model-theory and its ability to deliver language/world connections. The key point is that model-theoretic semantics may well be a theory which employs the R^0 notion of reference.

Other kinds of semantic theories more clearly explore some form of language/world relations, but they accomplish this in different ways. Let's distinguish reference in sense R^1 from reference in sense R^2, where R^1 takes the reference of a term to simply be some direct (perhaps causal) relation between that term and the world, and R^2 takes the relation to be rather more complex, involving at a minimum a four-place relation that involves the speaker, the expression used, context, and aspects of the world. One way of thinking about the difference is that on R^1 the term itself refers, but on R^2 a speaker uses the term to refer to certain aspects of the world under certain contextual conditions.

I gather that Chomsky takes R^1 to be what most direct reference theorists have in mind when they talk about reference, and indeed at least some of them may be using "reference" in that sense. On the other hand, it is clear that not all of them are (consider, for example, the literature on demonstratives, which routinely includes argument places for speaker and context). In any case, I propose that we set aside the R^1 notion of reference. For purposes of discussion, by "referential semantics" I will mean some form of semantic theory which employs R^2. I'll let individual theorists decide whether they subscribe to R^1 or R^2, although I suspect that many philosophers take R^2 to be the operative notion in their work.

Of course, saying that a semantic theory employing R^2 can express language/world connections still doesn't make matters crystal clear, since we still need to know what we mean by "language" not to mention what we mean by "aspects of the world." As we will see, for Chomsky everything turns on our understanding of these two notions.

For example, consider what turns on how we are to understand the fourth place of the R^2 reference relation. Clearly there are a number of ways to take the phrase "aspects of the world." We could take it as being about phenomenological aspects of the world, or about Humean sense data, or about the posited objects of science, or about objects in the world (in some robust sense of 'object'). We will come back to this issue a bit later, but suffice it to say, for now, that some ways of filling in the fourth place will result in semantic theories where Chomsky would take exception.

First, however, we need to take a closer look at the nature of language, and in particular at the conception of linguistic inquiry that Chomsky has articulated. A number of philosophers working in the theory of reference have conceptions of language as a social object – conceptions which are completely antithetical to

Chomsky's views on the matter. If we are not on the same page as regards the nature of language, then any sort of debate about language/world relations is going to be at cross purposes.

1.2 I-languages vs. E-languages

Here I follow Chomsky's (1986) distinction between I-language and E-language. Strictly speaking it may not even count as an actual *distinction*, since on Chomsky's view there is I-language, and then there is a collection of poorly defined if not incoherent views which might be lumped together under the general title of "E-language theories."

An I-language is not a spoken or written language (whatever they might be), but is rather a state of an internal system which is part of our biological endowment. I might have I-language representations of English sentences, but those internal representations are not to be confused with spoken or written English sentences. They are rather data structures in a kind of internal computational system that humans are born with and which they have co-opted for communication among other purposes.

Crucially, we have to understand the I-language computational system *individualistically*. That means that the properties of the system can be specified completely independently of the environment that the agent is embedded in. Thus, it involves properties like the agent's rest mass and genetic make-up (and unlike relational properties like the agent's weight and visibility to hungry tigers).

From the E-language perspective, on the other hand, a natural language is a kind of social object the structure of which is purported to be established by convention (however "convention" is to be understood), and persons may acquire varying degrees of competence in their knowledge and use of that social object. I gather that on Chomsky's views such objects would be of little scientific interest if they did exist, but that in any case such objects *don't* exist.

Consider a possible candidate for an E-language. The French Academy is supposed to dictate that French is to have a certain set of properties which cover the French lexicon, pronunciation, etc. The only problem is that the French Academy doesn't dictate as much as it thinks it does. At best, it dictates a small range of superficial rules about French, all the while relying upon shared tacit knowledge of I-language which provides the substrate upon which those prescribed rules are parasitic. What they have are a collection of proclamations about what the structure of French *should* be, which in total are not even sufficient to constitute a respectable candidate E-language.

In short, on Chomsky's view, E-languages (insofar as we can make sense of such things) would be unsuitable objects for naturalistic inquiry, but the point is academic, since there are no such things. I-languages are the only serious candidates for the objects of linguistic inquiry.

2 Are Referential Semantics for I-languages Possible?

Given the notion of I-language sketched above, we can take up the following question: is it possible to have a referential (R^2) semantics for an I-language?

Most would suppose that Chomsky's answer to this question is in the negative, and on the face of it, that answer is the natural one. An I-language, after all, is supposed to be part of an agent's psychology, and ultimately part of the agent's biology. The properties of an I-language are therefore (on Chomsky's view) individualistic – as noted above, they are properties that hold of the agent in isolation (like having particular genetic make-up).[4] The properties of a referential semantics, on the other hand, appear to be anything but individualistic. They are not properties that an agent can have in isolation, since they express relations between linguistic representations and, among other things, aspects of the world external to the agent.

But it seems to me that there is room to hold that so long as we avoid the R^1 notion of reference, and adhere to the R^2 notion, a theory which traffics in language/world relations is completely reasonable. How is this possible, given that I-language is individualistic and reference, even in sense R^2, apparently involves language/world connections?

The general phenomenon of individualistic and relational sciences informing each other is fairly common. For example, it is routine in primate ecology to draw upon facts about primate physiology (which is individualistic) to support a particular claim about the relation of the primate to its environment. More generally, as Webster and Webster (1988) observe, anatomical structure (which is presumably individualistic) can place constraints on the types of (relational) environmental functions that are possible, and vice versa.[5]

Chomsky has written little on this topic, but what he has written seems to be consistent with this general point of view. For example, in Chomsky 1995, we have the following:

> Naturalistic study is of course not limited to such [internalist] bounds; internalist inquiry into a planet or an ant does not preempt or preclude the study of the solar system or an ant community. Non-internalist studies of humans can take many forms: as phases in an Oxygen-to-Carbon Dioxide cycle or gene transmission, as farmers or gourmets, as participants in associations and communities, with their power structures, doctrinal systems, cultural practices, and so on. Internalist studies are commonly presupposed in others with broader range, but it should be obvious that the legitimacy of one or another kind of inquiry does not arise. (Chomsky 1995: 28)

The key thing to see here is Chomsky's view that internalist investigations can be presupposed by relational (i.e., non-internalist) investigations (although, at

least in this passage, Chomsky is conspicuously silent on whether the converse is true).

Of course, even if we allow that relational and individualistic sciences can inform each other, due to familiar arguments[6] one cannot draw *direct* inferences from the structure of an organism to organism/environment relations, but that does not mean that no inferences are possible. Rather it means that caution is necessary, and that each inference must be supported by careful argumentation and sound evidence – in short, it is the usual sort of scientific reasoning wherein nothing comes for free.[7] But the key point here is that our knowledge of the environment can inform our knowledge of I-language. Indeed, there is no need for Chomsky to take exception to the idea that knowledge of the environment might shed light on the nature of I-language.

So, then, what is all the noise about Chomsky and referential semantics? What precisely is the real bone of contention between Chomsky and most philosophers who espouse some form of referential semantics?

It seems to me that there are two deal-breakers here. One turns on how we are to understand the fourth place of the R^2 reference relation: "aspects of the world." As noted above, there are a number of ways to take this phrase (e.g., as being about phenomenological aspects of the world, or about Humean sense data, or about the posited objects of science, etc.).

There is another deal-breaker we should consider first, however, since it will help to clarify our discussion. This deal-breaker (a hypothesis that I will call language/world isomorphism (LWI)), strikes me as being at the root of many of Chomsky's concerns about referential semantics.

3 Language/World Isomorphism

The LWI label is a little bit crude, since the key idea is that there is an isomorphism holding between logical forms and the world. The basic idea is not restricted to the analytic tradition in philosophy. It has been at the root of the "semantic tradition" (in the sense of Coffa 1991) from Kant, through Bolzano and the early Wittgenstein, up to the present. As early as Kant, for example, we find some suggestion that there is an isomorphism between a representation and the thing represented.

> [Representation] is that determination of the spirit (*Bestimmung der Seele*) that refers to other things. What I call referring (*Beziehen*) is when its features conform to those of the external things.[8]

> [The representation] is composed out of its component concepts in the same way in which the entire represented thing is composed out of its parts. Just as, for example, one can say that the notes of a musical piece are a representation of the

> harmonic connection of the notes, not because each note is similar to each tone but because the notes are connected to each other just as the tones themselves.[9]

The most celebrated version of this general idea is surely the picture theory of meaning advanced by Wittgenstein in the *Tractatus*. On the view articulated there, just as "[w]hat a picture must have in common with reality, in order to be able to depict it – correctly or incorrectly – in the way it does, is its pictorial form" (Wittgenstein 1949: 2.17), a proposition must be isomorphic in structure to a state of affairs in order to represent or be about that state of affairs.

Of course we are not talking about the "surface form" of a proposition here.[10] We are talking about the proposition "under analysis" – in other words, we are talking about its logical form. For Wittgenstein, if one got down to the complete analysis of proposition it would reflect the structure of the world.[11] The key idea here is that we do not suppose there is an isomorphism between the surface form of a sentence and the world, but rather that there is an isomorphism between the sentence under analysis (i.e., its logical form) and the world.[12]

Most of the following arguments from Chomsky are often construed as being arguments against the possibility of semantics. In fact, however, it seems to me that they are not (by themselves) arguments against semantic theories which employ reference in sense R^2. They are rather arguments against any semantic theory which traffics in language/world isomorphism. Ultimately, they are intended to count against certain conceptions of referential semantics, but seeing how they work against these conceptions will require that we first carefully walk through some of these arguments.

4 Chomsky's Arguments against LWI

In Chomsky's recent writings there are a number of arguments which are designed to show that if we adopted a referential semantics (using R^1 or, I submit, any R^2 theory employing the LWI hypothesis), the kind of ontology we would predict does not appear to track our intuitions about the kinds of things we are really talking about. (i) Such theories allegedly commit us to things that we would never acknowledge as existing; (ii) they allegedly commit us to types of things which are different from the types of things that we ordinarily suppose we are talking about; and (iii) the things we talk about are of such a character that they are allegedly too unruly for such theories to handle. To have labels for all these possibilities, let's call them the "argument from implausible commitments," the "type mismatch argument," and the "misbehaving object argument."

I'll review all three of Chomsky's arguments before offering what I take to be the most natural reply to them. As we will see, this reply will lead us directly into the teeth of Chomsky's concerns about "aspects of the world," a topic that will occupy the final section of this paper.

4.1 The Implausible Commitments Argument

Chomsky draws attention to the fact that a referential semantics (apparently in conjunction with the LWI hypothesis) commits us to some apparently implausible entities, flaws for example, and suggests that we really can't take seriously a theory which commits us to such entities.

> If I say "the flaw in the argument is obvious, but it escaped John's attention," I am not committed to the absurd view that among things in the world are flaws, one of them in the argument in question. Nevertheless, the NP *the flaw in the argument* behaves in all relevant respects in the manner of the truly referential expression *the coat in the closet*. (Chomsky 1981: 324)

Pursuing a similar line of attack, Hornstein (1984: 58) has drawn attention to constructions like (1).

(1) The average man is concerned about his weight

Hornstein contends that "no one wishes to argue that there are objects that are average men in any meaningful sense."

The general concern introduced by examples like "flaw in the argument" and "the average man" is that the doctrine that there must be a tight connection between semantics and ontology is suspect, for it commits us to apparently absurd views.

This argument needs to be hedged up a bit before it is fully functional. As it stands, the argument makes the assumption that the logical form of these constructions is transparent – e.g., that the expression "the flaw" really is nothing more than a simple NP containing a determiner and a noun. This is not such an innocent assumption. After all, the LWI hypothesis does not hold that there is an isomorphism between surface linguistic form and the world, but rather between some ultimate logical form and the world. As Higginbotham (1985) has argued, before we admit a commitment to objects like flaws and average men, we would have to look more closely at the underlying logical form of these constructions.

With respect to the case of "flaw in the argument," Higginbotham suggests that the construction is parallel to "bad singer." The semantics of "bad singer" is not such that we say someone is a bad singer iff they are bad and are a singer. Rather, we understand that someone is a singer and that their singing is bad. Likewise we should not say that "that is a flaw in the argument" is true iff that is a flaw and that is in the argument. Rather, we should say that the phrase is true iff something is an argument and it is flawed. Despite appearances, "flaw" may not be a nominal, but may be a modifier like "bad."

Higginbotham offers a similar analysis for the case of "the average man." First,

he notes that the expression is actually ambiguous between an individual who has typical properties, and the sense in which we say that the average family has 2.3 children. It is presumably the latter sense which is problematic here. In this latter case, Higginbotham suggests that "average," despite appearing to be an adjective, is functioning as a kind of adverbial. The construction is parallel to examples like (2), discussed by Haïk (1983).

(2) Let's have a quick cup of coffee

Clearly in (2) we are not suggesting that there are cups of coffee which are in some sense quick. Rather "quick" is behaving as an adverbial, modifying the activity. Returning to "the average man," Higginbotham proposes that in a construction like (3)

(3) The average family has 2.3 children

the analysis will be something along the following lines:

(3′) On average, a family has 2.3 children

Chomsky might reply in a number of ways here.[13] Consider the rather more complex (4):

(4) Your report on the average family fails to make it clear that it has 2.3 children[14]

Even if suitable glosses can be worked out, they will be technical spellouts, of course, and these spellouts may introduce difficulties of their own. But the glosses themselves are not necessarily innocent. As Chomsky (pers. com.) asks, are we to continue with the strategy for examples like (5)?

(5) That his income is falling bothers John Doe

That is, are we to take John Doe to be standing proxy for "an average man," which in (5) is to be further unpacked with "average" operating as a kind of adverbial? This begins to constitute a very strong and possibly counter-intuitive thesis about the logical form of these constructions.

Returning to the case of "flaw in the argument," there are likewise responses to Higginbotham available. For example, Chomsky (pers. com.) asks how could one treat "flaw" as a kind of predicate in an example like (6)?

(6) We fixed three of the flaws you found but the rest of them resisted our efforts

One possible answer is to say that (6) can be glossed, as in (6′), where it is the steps in the argument that are intuitively flawed.

(6′) We fixed three of the flawed steps you found but the rest of them resisted our efforts

Again glosses of this nature have to be justified at some point, and there is a heavy burden to show that the introduction of this proposed hidden structure comports well with the rest of what we know about the syntactic form of these constructions. Still more, it has to be shown that general rules are available and not just case-by-case fixes. Any attempt to pursue the Higginbotham course here will not be a trivial exercise.

Even if the "average man" and "flaw" cases *can* be analyzed away, it is perhaps doubtful that *all* such cases can be analyzed away. One can, of course, imagine a project which would attempt to analyze away all "implausible" objects, but this would amount to rejecting certain core assumptions that we have agreed to grant Chomsky here. Such a project would essentially be committed to the super-externalist view that linguistic form could only be determined after ontology was complete – in other words, that linguistic form was determined by the embedding circumstances of the language. And that project, assuming it is coherent, would radically depart from Chomsky's individualistic assumptions about the nature of I-language.

4.2 The "Type Mismatch" Argument

Even if we are comfortable with perhaps counter-intuitive entities like flaws, Chomsky (1995) notes that there is an apparent mismatch between the type individuation that objects and substances intuitively have, and the type individuation that a referential semantics will provide. To get clear on the issue a terminological distinction will be helpful. Suppose we distinguish an I-substance from a P-substance, where a P-substance (if there is such a thing) is the sort of stuff that would play a role in physical theory (H_2O for example), and an I-substance (if there is such) is the stuff that we are intuitively talking about when we use language (the intuitive referent of "water"). To put the point in a more theory-neutral way, the I-substance is what it appears we are talking about based upon our use of language. The alleged problem is that P-substances and I-substances just don't match up right.

If, following Putnam 1975, "water" refers to H_2O, then a referential semantics will assign a P-substance (H_2O) as the semantic value of "water." But the problem is that the stuff we are actually talking about when we use the term "water" – the I-substance – is something else altogether. To see this, consider the fact that what we find in the Hudson River is called "water," though it could

hardly be considered H_2O. Also problematic is the fact that there are substances like ice tea which chemically approximate H_2O much more closely than Hudson River water, yet we don't call them "water." According to Chomsky, the situation is even more problematic than this. If someone at the water company poured tea leaves into the system so that what came out of the tap was chemically identical to Lipton Ice Tea, we would still call it "water" – although we might complain about its impurity.

So, here is the situation. What we are talking about when we use the term "water" – the I-substance – depends upon the social setting in which we find that substance. But according to referential semantics, the meaning of the term is supposed to depend upon the chemical composition of the substance referred to – it's supposed to be a P-substance. Conclusion: referential semantics (if respecting the LWI hypothesis) will not track the intuitive notion of meaning.

One might think it possible to get off the hook if one appeals to social theories of external (referential) content (in the sense of Burge 1979). Rather than P-substances, we might posit S-substances, substances that are individuated according to community norms. So, while my concept of water might not accord with H_2O, it might still accord with a certain socially determined object, which has the property of being water when it comes from the faucet, but not when it is served at a restaurant. The problem here is that there are plenty of examples where such S-substances (if there could be such things) would not track our intuitions about the extension of terms – Burge's own examples about arthritis and brisket are cases in point. That is, there is a mismatch between I-substances and S-substances. We will come back to this point a bit later.

4.3 The Misbehaving Object Argument

There are a number of interesting features to the water/tea story, one of which can be broken out as a separate objection to referential semantics.

We have already seen in the type mismatch argument that I-substances don't track P-substances, but there is another problem. The water/tea story also seems to show that the I-substance we are talking about when we use the term "water" is a most ill-behaved sort of substance. Something may cease to be water even if no internal physical changes have taken place. For example, the same chemical compound is water when it comes from the tap, but ceases to be water when it is served at a restaurant.

If that's the intuitive character of I-substances, then there is really little hope that referential semantics can "give the reference" of what we talk about when we talk about "water" and "tea," since referential semantics is supposedly going to say that the content of these terms is H_2O in the first case and H_2O plus certain other elements in the latter. That is, I-substances are so unruly that it is wildly implausible to suppose that they could have any counterparts in the

physical world. Hence they have no counterparts that a referential semantics could utilize as their referents.

A related, if somewhat more general, point is made in Chomsky (1975: 203), where Chomsky notes that the very notion of whether we are talking about a single object or a collection of objects turns on any number of social and institutional factors:

> We do not regard a herd of cattle as a physical object, but rather as a collection, though there would be no logical incoherence in the notion of a scattered object, as Quine, Goodman, and others have made clear. But even spatiotemporal contiguity does not suffice as a general condition. One wing of an airplane is an object, but its left half, though equally continuous, is not . . . Furthermore, scattered entities can be taken to be single physical objects under some conditions: consider a picket fence with breaks, or a Calder mobile. The latter is a "thing," whereas a collection of leaves on a tree is not. The reason, apparently, is that the mobile is created by an act of human will. If this is correct, then beliefs about human will and action and intention play a crucial role in determining even the most simple and elementary of concepts.

Moving that discussion into the current debate, we might say that it is implausible for even very simple semantic concepts like object and collection to correspond in any interesting sense with P-substances.

4.4 Analysis of the Arguments: Some Replies

The conclusion that we can draw from examples like the above is that any referential semantics purporting to respect the LWI hypothesis is going to misfire badly, since it is bound to utilize P-substances as referents, and P-substances just don't track the intuitive meanings of natural language expressions.

One reply for the semanticist is to simply reject the LWI hypothesis. That is, one could say that yes, a semantic theory is concerned about the connection between language and the world (as part of a four-place relation which also involves speaker and context), but this says nothing about there being an isomorphic mapping between primitive linguistic expressions (under analysis) and things (or kinds of things) in the world. If there is a single lexical entry corresponding to "water," it does not map onto any single substance, but rather various things depending upon circumstances, discourse participants, etc. Such theories are no less referential than those which respect LWI. After all, they are still in the business of articulating language/world relations, albeit through a very complicated story, and despite the fact that these relations are often dynamic one/many relations. I am guessing that Chomsky would have no objection to such a move.

There is, however, another maneuver – indeed, a standard philosophical

maneuver – which I gather is the place where Chomsky parts company from Putnam and Burge and the like. According to this strategy, individualist claims about the individuating conditions of "water" and "arthritis," etc. have no force. In the case of Burge's (1979) example of Oscar (who claims that he has arthritis in his thigh), one can say that "arthritis" refers (R^2) exclusively to ailments of the joints, and reject the claim that for Oscar, "arthritis" might apply to a painful condition of the thigh. Thus, we preserve R^2, LWI, and the intuition that the world is furnished by P-substances (or at least S-substances).

But now it is not clear that we can retain the I-language thesis. If there is an isomorphic relation between the content of Oscar's utterance and the syntactic form of that utterance (in other words, if LWI holds), and if the content of Oscar's utterance depends upon the environment in which he utters it, then it looks like the syntactic form of Oscar's utterances is going to depend upon the (possibly social) environment in which the utterance takes place. But then the syntactic form of Oscar's utterances doesn't depend solely upon facts about Oscar in isolation. And this is just to reject the I-language thesis.

Is it possible to retain I-language, LWI, and externalism about content? Could one argue that the linguistic representations remain unchanged but that their contents shift as Oscar moves between Earth and Twin-Earth? This seems unlikely. Following the dialectic of a related argument due to Boghossian (1997), consider the case where Earth and Twin-Earth differ only in that on Twin-Earth there is no gold (substitute metals are used). Twin-Earthians have the concept of gold, but it is by virtue of certain experts who have postulated the possibility of gold on the basis of physical theory. Thus, the concept of gold on Twin-Earth is a complex one, having to do with numerous facts involving the periodic table, etc. On Earth the concept of gold is an unstructured concept about gold itself. Now consider the cases of Oscar and Twin-Oscar, both of whom utter "I want some gold." Assuming isomorphism between linguistic structure and conceptual structure, when Oscar utters the word "gold," he employs a singular referring expression. Meanwhile, when Twin-Oscar utters "gold," he employs a structured description.[15] But by hypothesis Oscar and Twin-Oscar are internally identical. So, if LWI is respected, the difference in the syntax of their utterances must depend exclusively upon the respective environments in which they are embedded. Hence I-language has to be given up. We are forced to a notion of E-language, which Chomsky has already rejected as incoherent.[16]

Summarizing the strategies covered thus far, one can retain the notion of I-language with R^2 reference to P(S)-substances, but only if one gives up LWI. If we try to hang onto LWI and P(S)-substances, then we have to give up I-language. Is there another way out?

Why not give up the P(S)-substances? That is, why not bite the metaphysical bullet and acknowledge that there really are I-substances – that is, things like flaws, etc. They are clearly not logically absurd entities, and it need not be conceded that they are particularly odd entities. In the case of flaws, at least, one

might say that they are altogether common in the arguments one runs across, and one might wonder why they should be considered any less real than, say, tables and chairs. Likewise, coats hanging in the closet need not have any particular ontological priority over average families and flaws.

But is this a metaphysical bullet that we can bite? Can our ontology admit such things? Here (finally) we come to the issue of "aspects of the world" and the kind of things that we can be realists about.

5 Aspects of the World

Operating throughout all three of Chomsky's arguments above is the assumption (shared with Putnam et al.) that P-substances are the kinds of substances that a referential semantics is going to favor. For example, Chomsky (1992: 208) remarks:

> To be an Intentional Realist, it would seem, is about as plausible as being a Desk- or Sound-of-Language- or Cat- or Matter-Realist; not that there are no such things as desks, etc., but that in the domain where questions of realism arise in a serious way, in the context of the search for laws of nature, objects are not conceived from the peculiar perspectives provided by the concepts of common sense.

But why should we make the assumption that "the domain where questions of realism arise in a serious way" is "in the context of the search for laws of nature"? One possible answer would be that physical theory gets to say what's real, and hence if we are to have a genuine referential semantics in which the referents are "real" existing entities, then we are stuck with the kinds of entities and substances posited by physical theory. But this answer makes a strong assumption about scientific realism – one which is controversial to say the least, and most likely false, in my view.

There is a great deal of literature in the philosophy of science (for example van Fraassen 1980) which holds that the entities which science posits do not exist in the same sense as mid-sized, earth-bound objects like tables and chairs. Pursuing this line of thinking, we might say that scientific theories, despite their great interest and utility, are not the arbiters of what is real.

If we set aside the exclusive claim of the physical sciences on our ontology (i.e., if we dismiss P-substances as our semantic values), then we may well find that I-substances are entirely plausible candidates for the referents of a semantic theory. So, for example, it may be that the semantic value of "water" just is water – that complicated I-substance which moves through pipes into our homes, etc. If we take that route, then it is far from clear that the LWI hypothesis must be surrendered.

But perhaps the appeal to I-substances is a cheat.[17] Or perhaps it is entirely

parasitic on the notion of I-language representations. What is an I-substance if not simply "whatever corresponds to a particular I-language representation"? Seen in this light, isn't the talk of I-substance vacuous? Or at best, isn't it a misleading way of talking about I-language?

Here we need to head off a confusion. Even if it should turn out that I-language representations are in some sense prior to I-substances (perhaps logically prior or metaphysically prior) it does not follow that our path as empirical investigators will begin with the linguistic representations. The idea is that the linguistic representations will indeed underwrite our metaphysical intuitions, but that because of this we can expect our metaphysical intuitions to shed some light on the nature of I-language. To take a concrete example, the lexical entry for "water" when fully fleshed out is bound to be a very complicated representation (or at least to interact in complex ways with other I-language representations). Our knowledge of this fact does not come from direct investigation of the lexicon, but rather is guided by metaphysical intuitions about water like those that Chomsky evinces in his articles – intuitions which take into account rich contextual information about whether the material in question is coming from a faucet, or is being served at a restaurant.

It will not do to argue that Chomsky's water intuitions are not metaphysical – that they are only about his "I-concept" of water, or about the phenomenology of water. If there is no world beyond these concepts or beyond the phenomenology, then once we have made the step from talking about linguistic representations to talk of concepts, or phenomena, or sense data, we have basically stepped into the world – or as far into it as we are ever going to get.[18]

Of course, this picture is as Kantian as it can be. We have metaphysical intuitions, and we want to know what underwrites those intuitions. The first departure from Kant lies in the answer given – not "the categories of reason" but rather the structure of I-language. The second departure lies in the fact that our approach need not be entirely transcendental. We do have substantial independent knowledge of the language faculty, and we can use that knowledge to gain insight into the nature of reality.

6 Conclusion

I've made two claims in this chapter. The first one is that Chomsky is not hostile to referential (R^2) semantics *per se*, but rather toward certain auxiliary hypotheses about language/world isomorphism and toward ways we might understand talk of "the world." The second, somewhat contentious, claim is that language/world isomorphism is not implausible and that we can retain the hypothesis if we are prepared to be less ontologically stingy. More contentious still, the hypothesis may have some value to actual empirical research into the nature of I-language.

Naturally, caution is necessary in any attempt to draw conclusions about the nature of I-language from our knowledge of I-substances. Our knowledge of such substances is certainly fragmentary (are there really flaws and average guys or not?), and the nature of the link between those substances and their linguistic representations remains to be explored. One plausible hypothesis, however, is that any grasp we have on metaphysics is by virtue of our having the linguistic representations that we do. That is, were it not for our I-language representations having certain properties, we would find ourselves in a rather different world of I-substances. If this view should turn out to be true, then firm intuitions about I-substances can certainly shed light on the nature of I-language. Matters work in the reverse direction as well. It is entirely possible that if semantic theory requires that we quantify over flaws and average guys, then that helps shed light on our imperfect intuitions about ontology.

In sum, if referential semantics is possible, then it provides a new source of evidence (intuitions about I-substances) which could shed light on the nature of I-language. Likewise, the study of I-language could shed light on the nature of metaphysics. These possibilities illustrate the great promise of developing referential semantics for I-languages, and they also explain part of the motivation that some of us have for pursuing such theories.

Chomsky's reply: pp. 287–95.

Notes

1 This paper has benefited much from an extended correspondence with Noam Chomsky on these issues. Earlier versions of this material were presented to the Riejka Branch of the Section for Analytical Philosophy, Croatian Philosophical Society (1995), at the 2nd Central European Summer School in Generative Linguistics, Olomouc, the Czech Republic (1995), at the Department of Cognitive Science at Johns Hopkins University (1995), and at the Conference on Naturalizing Semantics, Maribor Slovenia (1996). I am indebted to Nenad Miscevic, Boran Bercic, Snjezana Prijic, and George Wilson for helpful discussion, to Noam Chomsky, Michael Devitt, Norbert Hornstein, and Richard Larson for comments on various drafts of this paper, and to Susan Larson for the references on kangaroo rats, etc.

2 So far as I know, Fodor may still hold some view like this.

3 It might be objected that we can always specify an intended model, and that the actual world should be the intended model. But Lepore (1983) has observed that this move will not work. If there is to be an intended model, then we should like to know how to identify it. But, on the face of it, this is no easy task.

> How much about a world do we need to know before we can distinguish it from all other worlds? Presumably a lot. There presumably is a class of worlds in which the number of trees in Canada is even and one in which the number is odd. So far are

> we from being able to single out the actual world from all others that we do not even know which class it falls in. But do we need to distinguish the actual world from all others to understand our language? (Lepore 1983: 184)

4 Here I am granting Chomsky the assumption that cognitive science generally is individualistic. This assumption, while plausible, is certainly not universally held. See Burge 1986 for the anti-individualism view, and Chomsky 1995 for a response.

5 One interesting example comes from the study of the anatomy of the kangaroo rat, where significant progress has been made by making judicious inferences from physical structure to organism/environment relations, and vice versa. When kangaroo rats were first discovered, biologists observed that they had particularly large middle-ear cavities. According to Webster and Webster (1988), biologists have posited a number of hypotheses about the function of the cavity, and the hypotheses based on anatomy alone appeared to be the weakest. For example Hatt (1932) argued that the enlarged cavities shifted the weight of the head to the posterior and thereby assisted upright saltatorial locomotion.

An alternative, and rather natural assumption, however, was that the enlarged cavities served the function of improved hearing. Research undertaken in the 1950s incorporated both anatomical and environmental studies and concluded that the enlarged cavity assisted in hearing, and subsequent work by Webster and Webster (1971) showed that the function was to pick up low frequency sound waves typically generated by predators like owls and snakes.

6 In popular literature, this point has been stressed numerous times by Stephen Jay Gould.

7 As Webster and Webster (1988) summarize the situation:

> Implicit in many morphological studies is the idea that structure determines, and therefore can reveal, function. Unfortunately this is an oversimplification. A more realistic view is that structure – whether of an entire organism or its parts – places constraints on what functions are possible (Gans, 1985) and may, to the observant, suggest some that are plausible.

8 Kant 1910–83: vol. 15, 76–7. Translation is from Coffa 1991: 31.

9 Kant 1910–83: vol. 15, 78. Translation from Coffa 1991: 31.

10 The very idea of a surface form is problematic in any case, since sentences (or utterances) do not come with their forms on their sleeves. The form that they receive when written, for example, is hardly a surface form in any interesting sense. Even word and sentence boundaries presuppose a fairly rich theory about the underlying form of our utterances.

11 As Wittgenstein explains the matter in his 1915 notebooks:

> It is clear that the constituents of our statements can and should be analyzed by means of definitions, and must be, if we want to approximate to the real structure of the statement . . . The analysed proposition mentions more than the unanalysed. The analysis makes the proposition more complicated than it was, but it cannot and must not make it more complicated than its reference was from the first. When the proposition is just as complex as its reference, then it is *completely* analysed. (Wittgenstein 1961: 46)

12 The basic idea and its application to questions of ontology can be illustrated by means of a brief detour through the history of the philosophy of language. In medieval philosophy of language (even earlier, actually), it was observed that standard assumptions about the connection between language and the world lead to certain surprising conclusions. Consider a simple sentence like (i):

(i) A dog barked

Specifically, if we regard "a dog" as a referring expression, then it appears that there must be some object to which it refers. But what is the referred-to object in this case? Some vague dog? Such proposals were entertained by medieval philosophers of language, and the consequences were taken quite seriously (and still are in, for example, Fine 1985). Clearly, one operative assumption was that if one successfully expresses a proposition by uttering (i), then the components of (i) must refer to something.

Much later, Meinong (1904) argued from the assumption that all expressions refer, to the conclusion that there must be nonexistent objects. So, for example, if all expressions refer and sentence (ii) is true,

(ii) Smith seeks the Golden Mountain

then "the Golden Mountain" must refer, not to an existing object, but to a *nonexistent* object.

Frege is often credited with showing us how to steer clear of the imbroglios created by sentences like (i). If we use the expression "some dog" we are not referring to some object or other, but rather are saying that some object in the domain is a dog. The logical form of (i) is best represented as in (i′), which says that there is an x which is a dog and which barked.

(i′) $(\exists x)(\text{dog}(x) \,\&\, \text{barked}(x))$

Here there is at most reference to the properties *doghood* and *barking*. If the semantics of (i) is as in (i′) then (i) is properly understood as being about these properties rather than as being about an abstract vague dog.

The same general strategy was exploited by Russell (1905) in his theory of descriptions. Against Meinong, Russell argued that definite descriptions like "The present King of France" do not refer, but rather *denote*. Thus, a sentence like (iii) has the underlying logical form in (iii′).

(iii) The present King of France is bald

(iii′) $\exists x(\text{Present-King-of-France}(x) \,\&\, \forall y(\text{Present-King-of-France}(y) \rightarrow x = y) \,\&\, \text{Bald}(x))$

(iii′) is not about any particular object according to Russell, but is rather about the property baldness and the property of being the present King of France. (iii′) asserts that the world is such that it contains one object which has the property of presently being the King of France, and everything which has that property is bald. (See Neale 1990 for discussion.)

The examples I have just given are cases where assumptions about the nature of reality have influenced our theory of the logical form of language. If we cannot countenance vague objects or nonexistent objects, then perhaps we need to take a closer look at the underlying form of our language. Alternatively, we can also find cases where beliefs about the underlying form of our language have led to substantive metaphysical claims.

For example, Davidson (1968/9) argued for the existence of events on the basis of the logical form of action sentences. To illustrate, we can easily make the inference from (iv) to (v):

(iv) John ate the chips gracefully

(v) John ate the chips

Davidson proposed that this inference was a logical inference, and that the inference could be made formally once the underlying logical form of "John ate the chips gracefully" was revealed. Davidson suggested that its logical form should be as in (vi):

(vi) $(\exists e)$ [ate(John, the chips, e) & graceful(e)]

Roughly, (vi) can be understood as saying that there was an event e which was an eating of the chips by John, and e was graceful. The conclusion (v) follows by simple conjunction reduction. Davidson took this to be evidence for the existence of events.

While Chomsky would certainly be sympathetic to a number of these technical proposals, I suspect he would *not* accept the notion that there are metaphysical consequences to these proposals. He would doubtless reject Meinong's assertion that "The Golden Mountain" must refer to some sort of object, and he would not feel compelled to give some alternative analysis of the expression *à la* Russell's theory of descriptions (the theory of descriptions might, however, be motivated by other considerations). But the reason for rejecting these metaphysical consequences has little to do with a rejection of referential semantics (in sense R^2). Rather, I think the metaphysical entailments are rejected because they rest on the assumption (apparently held by everyone from Russell to Meinong to the early Wittgenstein to Davidson) that there is an isomorphism holding between language (under analysis) and the world.

13 In the following discussion I am indebted to personal communication with Chomsky.

14 One possible response for Higginbotham here is to render (4) along the lines of (4′):

(4′) Your report fails to make it clear that on average a family has 2.3 children

There are, however, potential pitfalls in the (4′) gambit. For a start, the report described in (4′) need not be *about* the average family at all; it could very well be a report on milk consumption in urban areas. Is (4″) a possibility?

(4″) Your report on the average family fails to make it clear that on average a family has 2.3 children

Obviously not, since this brings us back to the apparent quantification over the average family. But perhaps the upstairs NP can be analyzed away in another fashion, along the lines of (4‴):

(4‴) Your report on what, on average, the state of families is, fails to make it clear that on average a family has 2.3 children

But this might not do either, since (4) does not say that the report is on the general *state* of families.

15 For more on the distinction between structured denoting expressions and unstructured referring expressions, see Neale 1993.

16 Chomsky has made numerous independent arguments against content externalism which I am passing over here, since my primary interest is in ferreting out semantical theories that are compatible with the notion of I-language. The philosophical consequences of content externalism are of course vast, and it may well be that it leads to consequences which, if taken to their logical conclusion, would prove to be reductios of the doctrine. For example, there are extremely complex issues about the consequences of the doctrine for the theories of self-knowledge, epistemic warrant, and memory. For a survey, see Ludlow and Martin (forthcoming).

17 Both Paul Horwich and Noam Chomsky have suggested as much to me.

18 Nor does it help to argue that the concepts mediate between language and the world. In that case, the same considerations apply to the concept/world relation. If conceptual structure underwrites our metaphysical intuitions about I-substances, then our intuitions about I-substances will provide insights into our conceptual structure, and our conceptual structure will provide insights into the structure of I-language. For the record, I consider conceptual structure to be a dispensable "middle man" here.

References

Boghossian, P. 1997: What the Externalist Can Know A Priori. *Proceedings of the Aristotelian Society*, 97, 161–75.

Burge, T. 1979: Individualism and the Mental. *Midwest Studies in Philosophy*, 4, Minneapolis: University of Minnesota Press.

Burge, T. 1986: Individualism and Psychology. *The Philosophical Review*, 95, 3–45.

Chomsky, N. 1975: *Reflections on Language*. New York: Pantheon.

Chomsky, N. 1981: *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, N. 1986: *Knowledge of Language*. New York: Praeger.

Chomsky, N. 1992: Explaining Language Use. *Philosophical Topics*, 20, 205–31.

Chomsky, N. 1995: Language and Nature. *Mind*, 104 (413), 1–61.

Coffa, J. A. 1991: *The Semantic Tradition from Kant to Carnap*. Cambridge: Cambridge University Press.

Davidson, D. 1980: The Logical Form of Action Sentences. In *Essays on Actions and Events*, Oxford: Oxford University Press.

Dowty, D., Wall, R., and Peters, S. 1981: *Introduction to Montague Semantics*. Dordrecht: D. Reidel.

Fine, K. 1985: *Reasoning with Arbitrary Objects*. Oxford: Blackwell.

Gans, C. 1985: Vertebrate Morphology: Tale of a Phoenix. *American Zoology*, 25, 689–94.

Haïk, I. 1983: Syncategormatic Adjectives. Manuscript. MIT, Cambridge, Mass.

Hatt, R. T. 1932: Vertebral Columns of Ricochetal Rodents. *Bulletin of the American Museum of Natural History*, 63, 599–738.

Higginbotham, J. 1985: On Semantics. *Linguistic Inquiry*, 16, 547–94.

Higginbotham, J. 1990: Contexts, Models, and Meanings: A Note on the Data of Semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, Cambridge: Cambridge University Press.

Hornstein, N. 1984: *Logic as Grammar*. Cambridge, Mass.: MIT Press.

Jackendoff, R. 1972: *Semantic Interpretation in Generative Grammar*. Cambridge, Mass.: MIT Press.

Jackendoff, R. 1983: *Semantics and Cognition*. Cambridge, Mass.: MIT Press.

Kant, I. 1910–83: *Gesammelte Schriften*, vols. 1–29. Edition by Königliche Preussische Akademie der Wissenschaften und Deutsche Akademie der Wissenschaften zu Berlin. Berlin: De Gruyter & Reimer.

Katz, J. and Fodor, J. A. 1963: The Structure of a Semantic Theory. *Language*, 39, 170–210.

Katz, J. and Postal, P. 1964: *An Integrated Theory of Linguistic Description*. Cambridge, Mass.: MIT Press.

Lepore, E. 1983: What Model Theoretic Semantics Cannot Do. *Synthèse*, 54, 167–87.

Levin, B. and Pinker, S. 1992: *Lexical and Conceptual Semantics*. Cambridge, Mass.: Blackwell.

Lewis, D. 1972: General Semantics. In D. Davidson and G. Harman (eds.), *Semantics of Natural Language*, Dordrecht: D. Reidel.

Ludlow, P. and Martin, N. (eds.) forthcoming: *Externalism and Self-Knowledge*. Stanford: CSLI Publications.

Meinong, A. 1904: Über Gegenstandstheorie. Translated as "The Theory of Objects," in R. Chisolm (ed.), *Realism and the Background of Phenomenology*, New York: Free Press, 1960.

Montague, R. 1974: *Formal Philosophy*. Ed. Richmond Thomason. New Haven: Yale University Press.

Neale, S. 1990: *Descriptions*. Cambridge, Mass.: MIT Press/Bradford Books.

Neale, S. 1993: Terms Limits. In J. Tomberlin (ed.), *Philosophical Perspectives*, 7: *The Philosophy of Language and Logic*, Atascadero, Calif.: Ridgeview Press.

Partee, B. 1980: Montague Grammar, Mental Representation, and Reality. In P. French, T. Euhling, and H. Wettstein (eds.), *Midwest Studies in Philosophy*, 5, Minneapolis: University of Minnesota Press.

Putnam, H. 1975: The Meaning of "Meaning." In K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, vol. 7: *Language, Mind and Knowledge*, Minneapolis: University of Minnesota Press.

Russell, B. 1905: On Denoting. *Mind*, 14, 479–93.

Van Fraassen, B. 1980: *The Scientific Image*. Oxford: Oxford University Press.

Webster, D. and Webster, M. 1971: Adaptive Value of Hearing and Vision in Kangaroo Rat Predator Avoidance. *Brain, Behavior and Evolution*, 4, 310–22.

Webster, D. and Webster, M. 1988: Hypotheses Derived from Morphological Data: When and How They Are Useful. *American Zoologist*, 28, 231–6.

Wittgenstein, L. 1949: *Tractatus Logico-Philosophicus*. London: Routledge & Kegan Paul.

Wittgenstein, L. 1961: *Notebooks 1914–1916*. Ed. and tr. G. E. M. Anscombe. New York: Harper & Row.