A Chomskian Alternative to Convention-Based Semantics

STEPHEN LAURENCE

In virtue of what do the utterances we make mean what they do? What facts about these signs, about us, and about our environment make it the case that they have the meanings they do? According to a tradition stemming from H.P. Grice through David Lewis and Stephen Schiffer it is in virtue of facts about conventions that we participate in as language users that our utterances mean what they do (see Grice 1957, Lewis 1969, 1983, Schiffer 1972, 1982). This view currently enjoys widespread acceptance among philosophers of mind and language. Though most are not particularly interested in the details of such programs, the dominant view seems to be that something of the sort proposed by Grice, Lewis and Schiffer is basically right. Thus, Jerry Fodor, reflecting what I take to be prevalent attitudes in the field, writes,

> [C]onsider the fact that tokens of "talcum powder tastes nasty" are true iff talcum powder tastes nasty. It may well be that the best story we can tell about that fact adverts essentially to certain communicative intentions of speaker/hearers of English. I don't offer anything like a detailed account of how this story might go. ... Perhaps it would implicate the speaker's intention that the token he produces should be *taken* to be true iff talcum powder tastes nasty; or perhaps what's crucial is the speaker's intention to adhere to a system of conventions, shared by members of his language community ... Though the details are disputed, some such account can be pieced together from the insights of philosophers like Grice, Schiffer, Lewis and Harnish, among others. (1990, p. 314)

I believe that convention-based accounts of the meaning of natural language utterances are fundamentally mistaken. They are accepted largely on the basis of the truism that language is in some sense conventional and because they are seen as "the only game in town". I do not dispute the fact that language is in *some* sense conventional; I do, however, dispute the claim that language is conventional in the sense that Lewis and other convention-based theorists have claimed. My discussion of convention-based accounts will focus on Lewis's version of the theory but I intend my argument to be general. I shall argue that convention-based accounts lack the generality we should want a theory of natural language semantic proper-

Mind, Vol. 105 . 418 . April 1996 **© Oxford University Press 1996**

ties to have, and more importantly, that there is a simpler, more general alternative to such accounts. My alternative is based on the standard, Chomskian interpretation of linguistics (see for example Chomsky 1965, 1975). This latter theory is not usually taken to be in competition with convention-based accounts of semantic properties. Indeed a central claim of my paper is that, contrary to popular opinion, the Chomskian account, suitably extended, should be thought of as being in direct competition with convention-based accounts. Once we see the Chomskian account as a rival to the convention-based account, the latter is no longer "the only game in town". And, I will argue, relative to the Chomskian account, convention-based accounts are unmotivated, lacking in generality and empirically unsupported.

The general structure of my argument is as follows. In the first section, I will argue that Lewis's technical notion of "conventionality" is not the sense in which it is truistic that language is conventional; thus Lewis's account is similarly not truistic. It is open to challenge. In § 2, I argue that if the concern is to account for the conventionality of language, Lewis's account should generalize, since syntactic and other linguistic properties of utterances are exactly as conventional (in the ordinary pretheoretic sense of "conventional") as are semantic properties. In § 3, I present my alternative to convention-based semantics and show that this account generalizes to other linguistic properties, in particular, that it provides parallel accounts of syntactic (and other linguistic) properties. I also present a number of empirical considerations in this section which favour the Chomskian account over the convention-based account. In § 4, I argue that looking beyond literal expression meaning does not help the case for convention-based semantics.

Before I start, though, I should note two disclaimers. First, though I call my alternative a "Chomskian" alternative to convention-based semantics, it is not necessarily one which Chomsky himself would approve of, not least because Chomsky himself is highly sceptical about the project of truth conditional semantics. I call it "Chomskian" because I take it to be broadly in the spirit of Chomsky's views about the nature of linguistics. I also want to be agnostic about the relation between the language processor and the grammar, so far as this is consistent with central assumptions of recent work in linguistics and psycholinguistics. For expository purposes, however, I adopt a relatively strong thesis that Chomsky himself would probably not accept, even for syntactic properties.

The second disclaimer is that in arguing against convention-based accounts of (expression) meaning I am arguing against a popular view that has its roots in Grice's work. I am not thereby arguing against all Gricean views about language. My account in effect rejects the primacy of speaker

meaning and the particular connection between contents of mental states and contents of utterances implicit in convention-based accounts. It does not (necessarily) thereby reject speaker meaning or communicative intentions, or beliefs regarding presuppositions of utterances, or beliefs arrived at by conversational implicature. As far as I can tell, my account leaves Grice's important work on conversational implicature in place, along with most other work in pragmatics. At the same time, I am not claiming that work on convention-based semantics stemming from Grice's 1957 paper "Meaning" should be interpreted as part of a theory of pragmatics. As we shall see, what I think is really at stake is how we construe semantics. In particular, the issue will be whether semantic properties should somehow or other be underwritten by our "general communicative abilities" or whether they should be taken to be more directly tied to features of the language processor, a special purpose cognitive mechanism for processing language.[1]

1. Two senses of "convention"

Lewis sees himself as defending the platitude that language is governed by convention. He uses the conventionality of language to pump our intuitions in favour of a convention-based account of the nature of semantic properties. His convention-based account, however, is based on a technical notion of "conventionality". In this section I present Lewis's account of "conventions" and the specific "conventions" which he takes to underwrite the semantic properties of natural language utterances. I argue that Lewis's account has powerful empirical consequences that are far from trivial or platitudinous—and that his account goes well beyond the intuitive pretheoretic sense in which language is thought to be conventional.

It may seem just *obvious* that language is conventional. After all, the semantic properties of natural language utterances are not *intrinsic* properties of the marks and sounds which have them. The noise "chocolate", for example, might not have meant *chocolate*. It might have meant *planet* or *train* or nothing at all. It might seem that this obvious conventionality lends a strong prima facie plausibility to the convention-based account.

[1] For some, this will in turn imply that the issue is whether semantics can be assimilated to the domain of the natural sciences. I think that if this assimilation is possible, then the enormously increased explanatory power it promises would provide a compelling reason for it. I further believe that it is possible to assimilate semantics to the domain of the natural sciences, and that the "Chomskian" account I advocate is the best strategy for accomplishing this.

David Lewis suggests as much in his motivational remarks at the beginning of his book *Convention*.[2] There Lewis says:

> It is a platitude that language is ruled by convention. Words might be used to mean almost anything; and we who use them have made them mean what they do because somehow, gradually and informally, we have come to an understanding that this is what we shall use them to mean. We could perfectly well use these words otherwise—or use different words, as men in foreign countries do. We might change our conventions if we like. (1969, p. 1)
>
> The platitude that there are conventions of language ... commands the immediate assent of any thoughtful person. (1969, p. 2)

We need *an account* of convention, though, in order to meet Quine's scepticism concerning the existence of such conventions of language. Lewis's project is to give an analysis of

> our common, established concept of convention, so that you will recognize that it explains what you must have had in mind when you said that language—like many other activities—is governed by conventions (1969, p. 3)

and thereby "rehabilitate analyticity" (1969, p. ix).

I completely agree that language is conventional in the trivial sense: it isn't an *intrinsic* property of the noise "chocolate" that it means *chocolate*. It might have meant *planet*, or whatever. I want to emphasize, however, just how big a step it is from this platitudinous sense of conventionality to the sorts of conventions Lewis ends up with and that convention-based accounts typically employ. It will make things easier if we introduce a small bit of terminology. I will refer to the intuitive pretheoretic sense in which language is thought to be conventional as "P-conventionality" (for Platitude-conventionality) and I will refer to the sense in which language is conventional on Lewis's account as "L-conventionality" (for Lewis-conventionality).

Lewis motivates his account of the conventions of language by noting several features we intuitively take to be part of the conventionality of language. The main one is that linguistic properties are not *intrinsic* properties of the signs which have them: it is in a certain sense arbitrary which

[2] And also in his later paper "Languages and Language": "It is a platitude—something only a philosopher would dream of denying that there are conventions of language, although we do not find it easy to say what those conventions are." (1983, p. 166)

linguistic properties are associated with a given utterance type.[3] As Lewis says:

> Words might be used to mean almost anything ... We could perfectly well use [our] words otherwise—or use different words, as men in foreign countries do. (1969, p. 1)

Since linguistic properties are only contingently associated with the signs which have these properties, they cannot be intrinsic properties of the signs.

Lewis also notes that though we might speak any of a large number of possible languages, we have a common interest in communicating, and therefore we have a common interest in speaking the same language (1969, p. 177). I agree with Lewis that these are platitudes. And I will take these two features together to characterize "P-conventionality".[4]

(1) *The P-conventionality of language*

1. Linguistic properties are not intrinsic properties of the utterances which have them. The marks and sounds we use could have (or could have had) linguistic properties other than those they in fact have.

2. Though we might speak any of a large number of possible languages, the members of a given linguistic community have a common interest in linguistic coordination (i.e., in speaking the same language), because they have a common interest in communication.

Lewis believes that his account unpacks the platitude that language is conventional and therefore provides a defence of the P-conventionality of language. He also intends his account to provide an explication of the nature of the semantic properties of natural language utterances (that is, an account of what makes it the case that natural language utterances mean what they do). In evaluating his account, then, we will have to keep both of these goals in mind. This section, however, is concerned only with the first of these goals.

[3] Utterance types are nonlinguistic physical types corresponding to some token utterance of a linguistic expression. So they are types of sounds (in spoken language), marks (in written language) or gestures (in sign language).

[4] In his motivational remarks, Lewis also suggests that it is part of the platitude that language is conventional that we

> have made [our words] mean what they do because somehow, gradually and informally, we have come to an understanding that this is what we shall use them to mean. (1969, p. 1)

It certainly seems true enough that we never all got together and formally decided at some point that we would use these words to mean what they do. And if this is all he means by this remark, I am willing enough to accept it. But Lewis seems to be making some stronger positive claim here, which he later spells out in his account of conventions (see below). I do not regard this further claim as a platitude about language, as I will argue below.

Lewis's account consists of a general account of conventions together with a specific convention governing language.

(2) *Lewis's General Account of Conventions*

> A regularity *R*, in action or in action and belief, is a *convention* in a population *P* if and only if, within *P*, the following six conditions hold. ...
>
> 1. Everyone conforms to *R*.
>
> 2. Everyone believes that the others conform to *R*.
>
> 3. This belief that the others conform to *R* gives everyone a good and decisive reason to conform to *R* himself. ...
>
> 4. There is a general preference for general conformity to *R* rather than slightly-less-than-general conformity—in particular, rather than conformity by all but any one. ...
>
> 5. *R* is not the only possible regularity meeting the last two conditions. ...
>
> 6. Finally the various facts listed in conditions 1 to 5 are matters of *common* (or *mutual*) *knowledge*: They are known to everyone, it is known to everyone that they are known to everyone, and so on. (1983, p. 164–6, numbering slightly altered from original.)

The specific convention involved in language use is a convention of truthfulness and trust in a given language. Truthfulness and trust are explained as follows.

(3) *Truthfulness and Trust*

> To be truthful in *L* is to act in a certain way: to try never to utter any sentences of *L* that are not true in *L*. Thus it is to avoid uttering any sentence of *L* unless one believes it to be true in *L*. To be trusting in *L* is to form beliefs in a certain way: to impute truthfulness in *L* to others, and thus to tend to respond to another's utterance of any sentence of *L* by coming to believe that the uttered sentence is true in *L*. (1983, p. 167)

Here a language *L* is a function (in the mathematical sense) from utterance types to sets of possible worlds, where an utterance of type *p* is true in *L* just in case the set of worlds which is the value of the function *L* for the argument *p* contains the actual world (1983, p. 163).[5]

The version of Lewis's general account of L-conventionality that we end up with for natural language, then, is the following (here I am just taking his account of "truthfulness and trust" and plugging it into his general account of conventions).

[5] This account of "true in *L*" gets significantly complicated in response to various facets of natural language such as indexicality, tense, ambiguity, mood and nonliteral uses of language (metaphor, irony, jokes, "white lies", etc.). Lewis says that he is "deliberately stretching the ordinary usage of 'true', 'truthfulness' and 'trust' in extending them" in these ways (1983, p. 172).

(4) *The L-conventionality of language*

1. There is a regularity of truthfulness and trust in *L* among the population *P*. This means that,

(*i*) for any given utterance, everyone tries to avoid producing that utterance, unless they believe it to be true in *L*,

(*ii*) everyone believes that the others do the same and

(*iii*) everyone responds to utterances by others by coming to believe that the uttered sentence is true in *L* because they believe the others try to avoid producing utterances unless they believe them to be true in *L*.

2. Everyone in *P* believes that the other members of *P* conform to the regularity of truthfulness and trust in *L*. This means that everyone believes that the others are doing the same as them regarding 1:

(*i*) for any given utterance they are trying to avoid producing it, unless they believe it to be true in *L*,

(*ii*) they are believing that the others are doing the same and

(*iii*) they are responding to utterances by others by coming to believe that the utterances are true in *L* because they believe that the others are trying to avoid producing utterances unless they believe them to be true in *L*.

3. The belief that the others conform to the regularity of truthfulness and trust in *L* gives everyone in *P* a good and decisive reason to conform to the regularity of truthfulness and trust in *L*. That is, believing that 2 holds provides us with good and decisive reason for doing 1.

4. There is a general preference for general conformity to the regularity of truthfulness and trust in *L*, rather than slightly-less-than-general conformity—in particular, rather than conformity by all but any one. Thus there is a general preference that everyone, rather than everyone but one,

(*i*) avoid producing any given utterance, unless they believe it to be true in *L*,

(*ii*) believe that the others do the same

(*iii*) respond to utterances by others by coming to believe that the uttered sentence is true in *L* because they believe the others try to avoid producing utterances unless they believe them to be true in *L*,

(*iv*) believe that the others believe others avoid producing any given utterance, unless they believe it to be true in *L*, and

(*v*) believe that the others respond by coming to believe that the uttered sentence is true in *L* because they believe the others avoid producing any given *L* utterance, unless they believe it to be true in *L*.

> 5. The regularity of truthfulness and trust in L is not the only possible regularity meeting the conditions 3 and 4. (So, for example, a regularity of truthfulness and trust in L' would meet the conditions 3 and 4 as well.)
>
> 6. "Finally the various facts listed in conditions 1 to 5 are matters of *common* (or *mutual*) *knowledge*: They are known to everyone, it is known to everyone that they are known to everyone, and so on." (1983, p. 166)

Certainly there is quite a difference between the claim that language is P-conventional and Lewis's claim that language is L-conventional. I particularly want to emphasize that Lewis's account has some powerful empirical consequences. I think this can be brought out in a variety of ways.

One way to see this is to notice that the account is committed to the existence of a wide range of propositional attitudes distributed amongst the population of speakers, in some cases propositional attitudes with quite complex and esoteric contents. So, for example, from 4.1 we all believe of our own utterances (i.e., of the marks, sounds or whatever) that they are true in L. For Lewis this means that we all believe that these utterances are in the domain of a function L mapping these utterances to sets of possible worlds which include the actual world. Actually this is a considerably simplified version of what we must believe since the meaning of the technical term "true in L" must be considerably complicated to deal with indexicality, tense, ambiguity, mood and nonliteral uses of meaning (metaphor, irony, jokes, "white lies", etc.)—all of these being commonplace in natural language use. Of course, we needn't believe that-our-utterances-are-mapped-by-the-function-L-to-sets-of-possible-worlds-which-include-the-actual-world, *under that very description*. But we do need to believe it under *some* description.[6] What's more, from 4.2 we must believe that others avoid making utterances unless they similarly believe this of their utterances. It is because we believe that they avoid making utterances unless they similarly believe that we ourselves respond to their utterances by coming to believe that p when they produce some utterance of $\ulcorner p \urcorner$ (that is, when they produce some utterance of the type x where $L(x)=p$). These beliefs explain why we ourselves conform to the regularity, according to 4.3. From 4.4, we must have a preference that everyone, rather than everyone but one,

(*i*) avoid producing any given utterance, unless they believe it to be true in L,

(*ii*) believe that the others do the same

(*iii*) respond to utterances by others by coming to believe that the uttered sentence is true in L because they believe the others try to

[6] See Schiffer (1987, pp. 258–61), for discussion of this point.

avoid producing utterances unless they believe them to be true in *L*,

(*iv*) believe that the others believe others avoid producing any given utterance, unless they believe it to be true in *L*, and

(*v*) believe that the others respond by coming to believe that the uttered sentence is true in *L* because they believe the others avoid producing any given utterance, unless they believe it to be true in *L*.

This is not to mention the fact that 1–5 are matters of common or mutual knowledge.[7]

A second way to see that the account is not simply platitudinous, but rather has strong empirical consequences, is to notice that it seems quite possible for the account to fail and yet for speakers to get on just fine. Imagine, for example, the following possible community. The community is small, unified and isolated and the members of the community are entirely unaware of the fact that natural languages vary. As far as they are concerned, there is only one language which everyone speaks. The speakers in this hypothetical community might well treat their linguistic symbols as if they had their linguistic properties *intrinsically*. It may never have occurred to them that the signs they use might not have meant what they do. When a speaker produces an utterance these speakers simply recognize that it has such and such linguistic properties. They needn't believe that the speaker believes the content. They needn't believe that the speaker believes that *the hearer* believes the speaker believes it. They needn't believe that everyone comes to believe that the utterance is true because they all know they all know that speakers try to avoid producing utterances unless they believe them true in *L*. And these beliefs certainly needn't provide them with any reason to conform to the regularity, or prefer that others do, since they needn't even recognize that alternative correlations of linguistic properties and sound types are possible. I see no reason why such a community should not be possible—indeed may well have been actual!

Finally, we can also see the empirical commitments of the account by noting some of the traditional objections which have been raised against it. For example, it is quite possible to produce utterances which you do not believe to be true in *L* (perhaps simply for fun) or which you do not

[7] Lewis qualifies this last point rather substantially in "Languages and Language", apparently in light of considerations of the plausibility of the "psychological reality" of the account:

> The knowledge mentioned here [in clause (6)] may be merely potential: knowledge that would be available if one bothered to think hard enough. ... Perhaps a negative version of (6) would do the job: no one disbelieves that (1) to (5) hold, no one believes that others disbelieve this, and so on. (1983, p. 165–6)

believe your audience will believe to be true in *L* (either because you have no audience or because your audience is extremely unlikely to be receptive to the message you have to convey).[8]

These points are only meant to establish that it is not a platitude that language is L-conventional. So while L-conventionality may provide an adequate explication of *conventions* (such as the convention of driving on the left side of the road in England)—as I am perfectly willing to grant for the purposes of this paper—it need not underwrite the nature of linguistic meaning. Once we see that it is not a platitude that language is conventional in the convention-based theorist's sense, we are free to consider challenges to the account seriously, and to explore alternatives to it. In § 2, I will develop a reason for thinking that the account should be challenged in this way.

2. *Conventions and other linguistic properties*

Lewis uses the P-conventionality of semantic properties to motivate his L-convention based theory of the nature of semantic properties, suggesting that his theory would defend the (P-)conventionality of language against philosophical attack. If Lewis's account were the correct account of the nature of natural language semantic properties then we could see how the semantic properties of natural language utterances would not be intrinsic properties of them. This follows pretty much directly from 4.5 above, which says, in effect, that other conventions (e.g., a convention of truthfulness and trust in some language other than *L*, say *L'*), might just as easily have obtained among the speakers of *L*. To explain the second feature of P-conventionality we don't really need to appeal to Lewis's account at all. Since this feature of the conventionality of language follows simply from the fact that if we didn't attribute more or less the same linguistic features to utterances we wouldn't be able to communicate using language. Since linguistic communication is in our interest, so is coordination.

So Lewis's account at least provides an account of the (P-)conventionality of the semantic properties of natural language utterances. The trouble is that semantic properties are not the only linguistic properties that are P-conventional. So are the syntactic, morphological and phonological features of language; and in just the same way. We can see this by considering some examples.

Consider the sound type corresponding to some utterance of (5)[9]

[8] For more on these traditional objections, see § 4 below.

[9] I will be referring back to numbered sentences in what follows. Such refer-

(5) Sue ate the chocolate.

The point about the semantic features of (5) being P-conventional is just that we might use (or might have used) the words in (5) to mean quite different things (or nothing at all). So, for example, we might have used (the sound type) "chocolate"[10] to mean *planet*, in which case (5) would mean something rather different from what it in fact means. It is in our common interest that all the speakers in our linguistic community mean the same thing by "chocolate", if we are to communicate using language. It would hamper communication if, by "chocolate" some meant *planet* and others meant *chocolate*, and communication would be impossible if this sort of variation were the rule rather than the exception. So it is in our common interest to use words to mean the same things.

However, much the same point can be made regarding the syntactic properties of an utterance. Clearly it is arbitrary what syntactic structure is associated with a given utterance type in just the way it is arbitrary what meaning is associated with it. Consider again sentence (5). Though it is the case that as we use (5) it has roughly the syntactic structure given in (6), we might use (or might have used) (5) in such a way that it had the syntactic structure given in (7) instead.

(6)

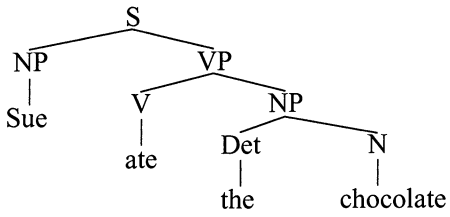

(7)

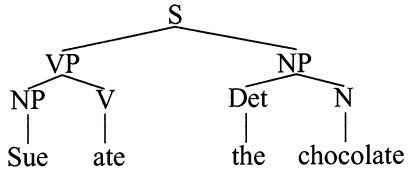

(6) has "Sue" as the subject, and "the chocolate" as the direct object, while (7) has "the chocolate" as subject and "Sue" as the direct object.

Clearly, it is also in our common interest that all the speakers in our linguistic community assign the same syntactic structures to expressions, if

---

ences should be interpreted as referring to the sound type corresponding to some utterance of that sentence (the same sound type throughout the discussion of that particular sentence) rather than to the sentence (type or token) itself.

[10] Taking this sound type to be a proper part of the one we are associating with (5).

we are to communicate using language. It would hamper communication if some assigned the structure (6) to (5) while others assigned the structure (7) to (5). Communication would be impossible if this sort of variation were the rule rather than the exception. So it is in our common interest to all use expressions so as to assign the same syntactic structures to them.

Similar points can be made regarding the conventionality of word boundaries or phonetic typing. We might use (or might have used) the noises that make up a spoken utterance so that word boundaries occurred at different places, and we might use (or might have used) the noises that make up a spoken utterance so that they were phonetically typed with different sets of noise types. So, for example (to illustrate the latter claim), the t sound in the word "tasty" (as it is used) in (some spoken utterance of) (8) might be typed together with all the p sounds, rather than the t sounds.

(8) Chocolate is tasty and nutritious.

Clearly, it would be in our common interest for all the speakers in our linguistic community to assign the same word boundaries to utterances and to type sounds the same way phonetically, if we are to communicate using language. It would hamper communication if, for example, some assigned to (8) the phonetic structure I would associate with (9) while others assigned to (8) the structure I would associate with (10).[11]

(9) Chocolate is tasty and nutritious

(10) Chocolate is pasty and nutritious

Communication would be impossible if this sort of variation were the rule rather than the exception. So it is in our common interest all to use expressions so as to assign the same word boundaries and phonetic structure to them.

Many (if not all) of the linguistic features of utterances are P-conventional. But how plausible is it that they are also L-conventional? How plausible is it that there are conventions (in Lewis's sense) which we participate in which make it the case that utterances have the phonological, morphological, or syntactic features they do? If Lewis is trying to defend the conventionality of language, the default position would be for his theory to generalize to cover these other linguistic properties. However, a convention-based account similar to that offered for the semantic properties of natural language utterances does not look at all promising as an account of these other sorts of linguistic properties. I don't believe that

[11] Differences of assignments of roughly this sort do exist among dialectical variants of a language; and, of course, these sometimes *do* lead to difficulties in communication!

anyone has ever suggested such an account of them.[12] This is most likely because no one seriously believes that speakers and hearers have the requisite attitudes or pattern of reasoning concerning the phonological, morphological or syntactic properties of utterances as is supposed in the case of semantic properties.

I think the moral is that the obvious conventionality of language gives us no reason at all to adopt a convention-based account of natural language semantic properties. Since syntactic, morphological and phonological properties are every bit as conventional in the trivial sense of conventionality (= non-intrinsic), but we are not in the least tempted to account for them in terms of the convention-based theorist's conventions, the fact that language is *obviously conventional* lends no credence to the convention-based account of semantic properties either. In fact, it gives us some reason doubt the convention-based account, since it fails to generalize to cover other properties which are equally obviously conventional. Fortunately, I think there is an alternative account which does not require us to make any use at all of Lewis's conventions to account for the various linguistic features (including semantic features) of utterances. I turn to that account now.

3. An alternative account

I noted at the outset that the convention-based account was the dominant view in philosophy concerning the nature of semantic properties of natural language utterances. Interestingly, the dominant view of the nature of the various other linguistic properties we have been discussing—syntactic, morphological, and phonological properties—is a very different view, derived from the work of Noam Chomsky. Since it seems to me desirable, if possible, to have a unified treatment of natural language linguistic properties (on general grounds of simplicity and theoretical elegance), and since we have already seen that the convention-based account does not seem promising in this regard, it seems worthwhile to explore the possibility of "expanding" the Chomskian account to cover semantic properties in addition to the various other sorts of linguistic properties. I think that this is possible, and in this section I argue that the Chomskian alternative

[12] I suppose that we could mimic the Lewisian conventions of truthfulness and trust, and attribute to speakers conventions of "grammaticality and trust", wherein speakers intended to utter phonologically, morphologically and syntactically well-formed utterances of their language, and people believed that others so intended and therefore interpreted the utterances of others as phonologically, morphologically and syntactically well-formed, and so on.

provides a satisfying account of the P-conventionality of language and provides a simpler and more general account of the nature of the linguistic properties of utterances than the convention-based account. I also note that the empirical evidence supports the empirical basis for the Chomskian view, but provides no equivalent support for the convention-based view.

Chomsky claimed that linguistics is a branch of cognitive psychology, that linguistic claims are claims about (a significant subcomponent of) our capacity to produce and understand our native language. According to Chomsky, our capacity to process language is not simply a reflection of our general cognitive capacity to reason. Rather, there is a special purpose cognitive mechanism responsible for this capacity (Chomsky 1965, 1975). The claim that there is a special purpose cognitive mechanism responsible for our ability to process language is now usually interpreted as the claim that there is a language processing module in the sense articulated by Jerry Fodor (1983): principally, a system that is domain specific and informationally encapsulated, whose operations are fast and mandatory. I take the claim that our linguistic abilities are in this sense modular to be empirically well supported. But I do not think that the Chomskian view I am considering necessarily requires this particular sense of modularity.[13]

Similarly, I take there to be a range of broadly "Chomskian" accounts of the nature of linguistic properties. I count a view as Chomskian if it treats the linguistic properties of utterances as inherited from features of the language processor. Chomsky himself explicitly says that he does not think that linguistics directly provides a theory of language processing, and he has had a somewhat sceptical outlook on developments in psycholinguistics. Still, Chomsky insisted that linguistic competence—what he takes linguistic theory to be a theory of directly—is a central and essential component of our language processor. I therefore take accounts of the nature of linguistic properties which link them essentially to features of the language processor to be broadly Chomskian in spirit.[14]

We may illustrate the view further by considering a specific "Chomskian" view in more detail.[15] According to this view utterances have the

[13] The issue of modularity in language processing is a focus of lively debate within psycholinguistics (see, for example, Marslen-Wilson and Tyler 1987, Frazier 1987, and other papers in Garfield 1987). Nonmodular parsers are not necessarily incompatible with the Chomskian account I present here since a parser could make use of contextual information and still successively compute the representations corresponding to the various linguistic levels of description. Furthermore, while nonmodular parsers would make use of contextual information, there is no reason to suppose that they would involve Lewisian reasoning.

[14] For discussion of some alternatives see Stabler (1983) and Matthews (1991).

[15] The particular account I offer is not the only possible Chomskian account. The general model of language processing and the suggestion I present as "the Chomskian alternative" are offered for the sake of concreteness and expository

linguistic properties they do in virtue of being associated, in the course of language processing, with mental representations having those properties.[16] To see how this account might work, we can adopt the following general picture of language processing. A sentence, (8) for example, is uttered.

(8) Chocolate is tasty and nutritious.

Processing is accomplished by successive computation of representations of this sentence at the phonological, syntactic and semantic levels. The processing proceeds according to general principles and on the basis of information about the linguistic properties of words stored in the mental lexicon and the input representations from the prior level of processing (see Forster 1979, Garrett 1990, Frazier 1988). Such a model is an idealization. It is not really true, for example, that the processor computes the complete syntactic representation of a sentence before passing that representation on to the semantic level. Still, the majority of linguists and psycholinguists working within Chomsky's broad theoretical framework would accept this picture in rough outline.[17]

According to this picture though, the processor constructs representations corresponding to the linguistically significant levels of description, including a semantic level of processing, corresponding perhaps to the linguistic level of "logical form", which, following Robert May, we may think of as

---

simplicity. As noted above, Chomsky's suggestion points us towards a certain range of accounts about the nature of linguistic properties, each compatible with a range of general processing models. A different and perhaps more plausible model might posit a less direct relation between the utterance and the linguistic properties, perhaps in terms of a modal relation of some sort (say, constructability given the nature of items stored in a mental lexicon and combinatorial rules represented in a stored mental grammar). Which sort of model is best will depend on various empirical and theoretical considerations (about which not enough is now known to make any clear choice). I take the detailed nature of the relationship between linguistic theory (semantics included) and the psychological mechanisms underlying our capacity to process our native language to be a matter for empirical psycholinguistic research. This does not mean that the view is vacuous or impossible to refute. It may well turn out that none of the claims in this general range can be sustained.

[16] See J.A. Fodor (1975) and Matthews (1991).

[17] Not all psycholinguists would accept this model, even in rough outline. Not all psycholinguists, for example, would even accept that there are distinct levels of processing corresponding to the linguistically significant levels of description (see, e.g., Marslen-Wilson and Tyler 1987). If these theorists are correct, then the specific Chomskian theory I present cannot be right (though of course some other broadly Chomskian account might be—see note 15). Naturally I take it to be an open empirical question whether the empirical consequences of the actual Chomskian account I present obtain.

> that level of representation which interfaces the theories of linguistic form and interpretation. ... it represents whatever properties of syntactic form are relevant to semantic interpretation—those aspects of semantic structure which are expressed syntactically. (1985, p. 2)

On this version of the Chomskian view, the semantic properties of utterances would be thought of as being "inherited" from the semantic properties of the representations at this level, and, in general, the linguistic properties of utterances would be inherited from the associated representations at each of the various levels of processing. The model I have in mind here is actually very straightforward. Given the empirical claim that language processing consists in recovering a series of representations at various linguistic levels, the view is simply that *it is in virtue of being associated, in language processing, with these representations that an utterance has the linguistic properties it has*. So, just as an utterance has a certain syntactic structure in virtue of being associated with a representation which has that structure, so it has a certain content or meaning in virtue of being associated with a representation which has that content or meaning.[18]

This leaves unexplained in virtue of what an internal representation could have semantic properties in the first place. But this is not an explanatory debt peculiar to the Chomskian theory. Lewis's account also appealed to internal states with unanalysed semantic content (e.g., beliefs). In both cases the project is simply to reduce the problem of providing an account of natural language semantic properties to the problem of providing an account of the semantic properties of mental states. Both accounts simply assume internal states with semantic content. The difference is simply a matter of which sorts of internal states with content are supposed to underwrite the natural language semantic properties.

That, in rough outline, is the Chomskian alternative. Like Lewis's account, it is compatible with and can account for the fact that semantic properties of utterances are P-conventional. The first and main feature of the P-conventionality of language (1.1 above) is that linguistic properties are not intrinsic properties of the utterances that have them. That means that the marks and sounds we use could have (or could have had) linguistic properties other than those they in fact have. On the Chomskian account under consideration, linguistic properties are not intrinsic in that the mechanisms responsible for processing might have been constructed dur-

[18] A different, related, account would have it that the representations constructed in language processing were representations to the effect that such-and-such an utterance had so-and-so linguistic properties. Since my aim here is not to decide which amongst various competing Chomskian accounts is best, I prefer the simpler account in the text for expository purposes. Again see note 15 above.

ing language acquisition so as to respond to different utterances the way they now respond to English utterances, and they might have been constructed so as to respond differently (or not at all) to English utterances. Similarly, these mechanisms might be altered in such a way as to respond differently to English utterances, or to respond to different utterances as they now respond to English utterances. Limits are set on what sorts of variation are possible in attaching linguistic properties to utterances by our capacity for language acquisition and the malleability of the existing processing mechanisms.

The other feature of P-conventionality is that it is mutually in our best interest to coordinate our language use. As was noted above, no theory-specific considerations are required to explain this feature of P-conventionality. The account of this feature is just the same as that given above: if we did not attribute more or less the same linguistic features to utterances we would not be able to communicate using language, and given that linguistic communication is in our interest, so is coordination.

Unlike Lewis's account, the Chomskian account I have presented also generalizes to other linguistic properties. Moreover, it also provides an account of the nature of linguistic properties generally—providing equally good accounts of the *nature* of the phonological and syntactic properties of utterances as it provides of the nature of their semantic properties. The Chomskian account is also simpler, or more direct. No complicated set of beliefs or intentions is required: we just look to the mental representations directly associated with the utterance by the language processor.

The crucial point, however, is a more directly empirical point: the language processing mechanism does not need to make any use of the mental states Lewis's theory posits. If language processing is accomplished by a special purpose cognitive mechanism, there is no need to reason along Lewisian lines in order to process language or recover the semantic properties of utterances. The language processor can be thought of as treating utterances *as if* they had their linguistic properties *intrinsically*, in the sense that *the processor does not take the arbitrariness of linguistic properties into account at all in processing utterances*. In assigning syntactic structure to strings, for example, the processor seems to take into account only information stored in the mental lexicon about the words in the utterance such as their syntactic category (noun, verb, ...) and subcategorization information (e.g., that "gave" takes a direct object and an indirect object), and general principles of assigning structure such as Minimal Attachment or respect for the Theta Criterion. The language processor does not need to reason along Lewisian lines because it has access to special linguistic "knowledge" and processing principles that do all the work.

It is easier to see now how the hypothetical community considered in § 1 above is possible. Indeed, in many respects, we are such a community, since though it is possible for us to recognize the conventionality of language, by and large this recognition plays no role in our linguistic dealings with one another.

In addition to the automaticity of language processing, a variety of empirical considerations can be marshalled both in favour of the Chomskian account and against the convention-based account. I want to focus my attention here on the latter sort of considerations.[19] I take the real *core* of the convention-based theory to be that linguistic communication is to be assimilated to communication in general. Lewis, like other convention-based theorists, completely embeds his discussion of linguistic communication into a larger discussion of communication generally, and clearly emphasizes the affinities between the linguistic and nonlinguistic cases. In linguistic communication, as in communication generally, we engage in a process of reasoning to determine what the other person is trying to communicate. I point to me and then to a distant spot and point to you, and then the same spot, and you infer by some process of reasoning that I want to meet you over there. Linguistic communication is much the same, according to this view. It is this core, and not the details of any specific account, that I want to look at now.

Since language, on this view, is a rather direct reflection of general abilities to reason about one's own mental states and the mental states of one's conspecifics, we should expect the ability to use language to correlate strongly with these general capacities. In particular, given the obvious advantages of having a language, we should expect any agent with these basic capacities to develop linguistic abilities. Given that linguistic communication is simply one facet of a more general communicative capacity, and essentially involves some sort of reasoning about communicative intentions, we should expect that linguistic ability would strongly correlate with both communicative ability and general intelligence and reasoning ability. A number of considerations, however, suggest that this correlation does not in fact hold. I will break the argument down into several stages, each of which presents independent considerations, but considerations which I think are mutually reinforcing.

First, general intelligence does not correlate with linguistic ability. The correlation breaks down in both directions. There are severely intellectu-

[19] Among the forms of evidence of the former sort are various more direct types of psycholinguistic evidence for levels of processing from speech errors (Garrett 1988), from processing breakdown in normal and aphasiac subjects (Pritchett 1992, Zurif 1990), and from experiments involving cross modal priming (Swinney et al. 1988, cited in J.D. Fodor, 1990), to cite just a few of many possible sources of evidence.

ally challenged agents with normal linguistic abilities and there are intelligent agents who do not possess anything approaching normal linguistic abilities. The former sort of case is illustrated by subjects with Williams syndrome. Steven Pinker comments on a number of such cases in *The Language Instinct* (1994). To illustrate the linguistic abilities of such subjects, consider the following passage taken from a transcript of a woman named Crystal who has Williams syndrome.

> This is a story about chocolates. Once upon a time, in Chocolate World there used to be a Chocolate Princess. She was such a yummy princess. She was on her chocolate throne and then some chocolate man came to see her. And the man bowed to her and he said these words to her. The man said to her, "Please, Princess Chocolate. I want you to see how I do my work. And it's hot outside in Chocolate World, and you might melt to the ground like melted butter. And if the sun changes to a different colour, then the Chocolate World—and you—won't melt. You can be saved if the sun changes to a different colour. And if it doesn't change to a different colour, you and Chocolate World are doomed. (Pinker 1994, p. 53)

Clearly this woman has intact linguistic abilities. But, according to Pinker, Crystal and other people with Williams syndrome have an IQ of about 50 and are not able to do such simple things as find their way home, add, or retrieve things from the cupboard (Pinker 1994, p. 52).

The same point can be made even more dramatically by considering the case of a subject, Christopher, studied by Neil Smith and Ianthi-Maria Tsimpli (1991 and 1995). According to Smith and Tsimpli, Christopher has a nonverbal IQ between 60 and 70 and finds ordinary tasks like "doing up a button, cutting his fingernails or vacuuming the carpet" to be "tasks of major difficulty" (Smith and Tsimpli, 1991, p. 117). In spite of this his English is "entirely comparable to that of normal speakers" (1995, p. 44) and, amazingly, "when given a passage written in any of some 15 or 16 languages—[he] simply translates it into English at about the speed one would normally read aloud a piece written in English" (1991, p. 317). Though his competence is rather varied in the different languages, his overall abilities are impressive on any scale. Here are two examples from Smith and Tsimpli (1991, pp. 319–20):

> *Greek* (the passage was in Greek script)
> Otan perase t'amaksi, epsakse ja tis pantufles tis, ala ena paljopedho ihe pari ti mja ki efevje jelontas.
>
> *translation*: When the car passed, she looked for her slippers, but a naughty (lit: old) child had taken one and was leaving laughing.
>
> *C's translation*: "When she passed the car ... when the car passed, she was looking for her slippers, but an old child had taken one away and left ... and was laughing."

> *Polish*
> Musiałem go wrzucić do wozu siłą. Położył się na pldłodze i zamknął oczy, nie chcąc widzieć, co go jeszcze czeka.
>
> *translation*: I had to throw him into the car with force. He lay down on the floor and closed his eyes, not wishing to see what awaited him.
>
> *C's translation*: "I had to take him out of the car strongly and put—he put himself on the floor and opened his eyes—and shut his eyes, not wishing to see what was waiting for him."

Smith and Tsimpli (1991, pp. 318–22) also provide examples of his translations of passages from Danish, Dutch, Finnish, French, German, Hindi, Italian, Norwegian, Portuguese, Russian, Spanish, Swedish, Turkish and Welsh.

On the other hand, there are also cases of intelligent agents who do not possess anything approaching normal linguistic abilities. The classic cases are those involving persons exposed to language after the critical period for language acquisition. (We might also note that some apes have IQs as high as 80 (Wallman, 1992, p. 20), but language skills not even approaching Crystal's much less Christopher's!) Perhaps the most famous case involves a woman named Genie who was the victim of severe abuse and neglect as a child. Until the age of 13 she grew up in isolation and was almost never spoken to. Needless to say, when she was discovered she had very little cognitive or linguistic abilities, and was severely emotionally damaged. It turned out that though she was able to recover significantly, she was never able to acquire anything like normal linguistic abilities. Susan Curtiss, who was one of the main researchers to study Genie's development, describes her as a "powerfully effective nonlinguistic communicator" (1988, p. 98) but notes that her knowledge of the basic rules of English did not develop past that of a 2 year old in the 8 years in which she was studied. Here are some illustrative examples of Genie's speech:

> "applesauce buy store"
> "man motorcycle have"
> "Genie bad cold live father house" (Curtiss, 1988, pp. 98)

Genie's case is difficult because of the severe abuse she suffered as a child. The deficits may be explainable in other ways. The case of another woman, Chelsea, is perhaps better in this respect. Chelsea is a severely hearing impaired person who was misdiagnosed as being mentally retarded as a child. Since people attributed her inability to learn a language as being due to her alleged mental retardation, rather than her inability to hear the language, no effort was made to expose her to a sign language. When it was discovered that she had a severe hearing impairment (unbelievably, in her 30's), she was fitted with hearing aids, and an attempt was made to teach her language. Like Genie, she has not been able

to acquire normal linguistic abilities, despite being otherwise of normal intelligence. Here are some examples of Chelsea's speech:

> "The small a the hat"
> "Breakfast eating girl"
> "They are is car in the Tim" (Curtiss, 1988, p. 99)

Unlike Genie, she makes free use of determiners, prepositions and such. But her use of language is often so ungrammatical that it is unintelligible.

Cases like those of Crystal, Christopher, Genie and Chelsea are typically taken to show that linguistic ability is in some sense modular, since it can be dissociated in these ways from general intellectual ability. Possibly, however, general intelligence is not to the point in the case of convention-based semantics. The core of the convention-based account is an assimilation to general communicative abilities, not general intelligence. However it is worth noting just how severe this dissociation can be. Christopher and Crystal, we are told, find such tasks as buttoning a button or adding to be extremely difficult, and yet they are clearly rather proficient language users. Yet according to the convention-based account, language users must engage in some (notoriously) rather complex reasoning processes, at times involving quite complex and esoteric contents.[20]

The next stage in the argument is based on the fact that similar dissociations can be found concerning the possession of a so-called "theory of mind"—a general ability for reasoning about the mental states of others in terms of, for example, propositional attitudes. Evidence from recent studies of autistic individuals suggests that a central component of autism is the lack of a theory of mind (for a review of some of this literature, see Happe 1994). Although most autistic individuals have poor or nonexistent linguistic abilities and very low IQs, some "high functioning" individuals have normal IQs and, despite some rather serious communicative abnormalities (being withdrawn, or overly inquisitive, or otherwise socially inappropriate), they can also have nearly normal linguistic abilities. Genie and Chelsea illustrate the other side of the dissociation here as well, apparently possessing normal theory of mind abilities, but lacking linguistic abilities. This argument poses a more direct challenge to the convention-based theorist. It is extremely puzzling how someone lacking a theory of mind could use language according to the convention-based theorist, given that the account posits numerous attitudes concerning the propositional attitudes of other members of the community. It is also extremely puzzling why someone possessing a theory of mind and normal intelligence should fail to acquire normal linguistic abilities, given the

[20] The Chomskian explanation for their linguistic abilities is that they are compatible with severe nonlinguistic deficits because linguistic ability has little or nothing to do with "general intellectual abilities".

obvious advantages of being able effectively to communicate using language.

The next stage in the argument is based on evidence for the dissociation of general linguistic ability and the capacity for communication. One source of evidence comes from aphasiacs, where much of the ability to use language may be lost (though not through motor damage), but general communicative skills seem to remain intact. A particularly interesting case here involves aphasiac "speakers" of American Sign Language (ASL), since in this case linguistic and nonlinguistic communication is often in the same (gestural) modality. Poizner, Bellugi and Klima (1987) tested a number of such aphasiacs for their ability to imitate nonlinguistic representational gestures, to imitate nonlinguistic *non*representational gestures, and to interpret pantomime. All but one[21] performed normally on these tests, though each had severe linguistic deficits.

An even more striking case arguing much the same point is provided by the case of children learning ASL. When English speaking children first learn the pronouns "me" and "you" they often make errors, referring to themselves as "you" and to others as "me". In ASL, the signs for "me" and "you" are pointing gestures—pointing to me for "me" and you for "you". Deaf children, like hearing children regularly point to objects of interest in their environments in their prelinguistic phase. So one might expect that children learning ASL would not be prone to the sort of reversal errors that children learning English pronouns are, given that the ASL sign, and the prelinguistic sign, are physically identical, and the children seem to have mastered the sign in their prelinguistic phase. We would be especially likely to suppose this, of course, if we thought that natural language was basically just a reflection of a general communicative ability, as the convention-based account suggests.

It turns out, quite surprisingly, that children learning ASL make the same sorts of reversal errors as children learning English. Children learning ASL stop using pointing gestures to refer to people for a while around the age of two, and then shortly after, at about the same age as children learning English acquire pronouns, pointing gestures reappear—with the reversal errors (Petitto, 1987).

If language were a direct reflection of general abilities to reason about one's own mental states and the mental states of others, we should expect the ability to use language to correlate strongly with these general communicative capacities. We would not expect radical discontinuities between linguistic and nonlinguistic use of what is physically *the same*

[21] The one exception was to the test of imitating nonlinguistic representational gestures, and the deficit was predictable from the particular lesion involved, according to current theory: Poizner et al. 1987, pp. 170–2.

symbol. The fact that it does not correlate and that such discontinuities exist suggests that the ability to speak a language is not simply a reflection of general communicative or intellectual ability.

The final sort of evidence I will consider shows how specific linguistic constraints (which appear to be due to Universal Grammar) can override what would otherwise seem to be decisive general communicative considerations—namely, fitting the language you "learn" to the model you are given. One striking sort of case involves children who actually *opt out* of the "linguistic" practices in their "linguistic" environment. Deaf isolates are deaf children born to hearing parents and not exposed to any natural sign language. These children construct for themselves a sign language with many of the marks of standard natural languages, despite the fact that the only "language" they are exposed to is the primitive signing of their parents, which does not constitute a natural language.[22] Amazingly, the linguistic developmental pattern of these children is much the same as normal children (through age 3 or 4) despite the fact that they effectively have *no model at all*, and therefore *no training* in the use of this language (since they are making it up!). The children's language takes on a number of language-like features that go beyond the parental input. For example, the children use pointing gestures to indicate not just objects in the immediate environment, but also objects located in imaginary space, not present in the immediate environment, and at least some of these children develop their own inflectional system for verbs where action signs are displaced toward the position of the sign for the object in motion (a common device in standard sign languages).[23] Moreover, the parents' sign system never develops to the same degree as the children's, and when the parents' system does develop to some degree, it is in *response* to the spontaneous changes in the children's language and not the other way around. Also, the children's "vocabulary" overlap with their mothers' symbols is only an estimated 33% (Jackendoff 1994, p. 129).

Exactly the same sort of phenomenon seems to occur when children are exposed to a pidgin language: they seem to opt out of the "linguistic" practices of their "linguistic" community. A pidgin is a makeshift amal-

[22] Though I take it to be a theoretical issue (to be settled by linguistic theory) just what counts as a natural language, we can note here briefly that the "language" which these children's parents use to communicate with them typically lacks many important features of natural languages (for example, they typically have no inflectional system). For discussion, see Jackendoff 1994, pp. 126–30 (on which my remarks are based), or the original research Jackendoff describes there (Goldin-Meadow and Mylander 1990 and Goldin-Meadow et. al. 1994).

[23] Data is rather limited due to the huge effort required to interpret the children's language with strict controls to ensure maximally objective coding. For more detailed information, see Jackendoff 1994, pp. 126–30 and the papers he cites.

gam of several natural languages used for communication among a group of speakers with no single dominant language. Children exposed only to pidgins reject them in favour of new languages of their own creation (creoles). The creole which children create is far richer and more systematic than the pidgin on which it is "based" and is uniform across the community (see Pinker 1994, Ch. 2 and Jackendoff 1994, pp. 130–5). These children effectively turn the primitive signing of their parents into a natural language through expansion and regularization of various sorts. What is interesting from the current perspective is that in going beyond the "language" of their environment, they are ignoring the conventional assignments of linguistic properties to utterances which are found in their "linguistic" environment: thus they would be opting out of the conventions in their linguistic environment. On the standard convention-based account it is hard to see why they would do this.

These phenomena ranging from the cases of Williams syndrome, to autistic individuals, to individuals past the critical period, to the deaf isolates and children exposed to pidgins pose very serious empirical difficulties for the convention-based account (while lending considerable support to the Chomskian account). In my view, they leave no serious doubt that the convention-based account is fundamentally mistaken.

4. Looking beyond literal expression meaning

In § 3 we saw that standard accounts treat language processing as a more or less automatic and autonomous process. Such an account leaves no real role for the various propositional attitudes invoked by the convention-based account. Further empirical considerations suggested that, contrary to the picture of linguistic communication implicit in such accounts, linguistic communication is not continuous with nonlinguistic communication or a general capacity for reasoning and intelligence.

We are not, however, forced to accept the view that speakers have no propositional attitudes about themselves, their partners in communication, or their mode of communication, nor need we claim that such attitudes play no role at all in linguistic communication. The model of language processing in § 3 is compatible with the existence of a level of processing at which language users employ general cognitive faculties in attempting to infer speakers' beliefs, intentions and expectations, and we have powerful reasons to suppose that there is such a level of processing. If we did not monitor the flow of discourse in this way, how would we be able to infer the nonliteral meanings of utterances or know what sort of response was appropriate in a given circumstance? I will call the "level"

at which we keep track of the beliefs and intentions of our conversational partners (as well as various other information pertaining to the immediate environment, and relevant general world knowledge) the "level" of discourse monitoring.[24]

Several possibilities for the convention-based theorist present themselves at this point.[25] Perhaps this is where the Lewisian reasoning that determines the literal semantic properties of utterances takes place. Or perhaps it is not literal semantic properties of utterances that we get an account of, but some other level of meaning which includes the nonliteral meaning of utterances, and perhaps the Lewisian reasoning which determines *these* properties of utterances takes place at the level of discourse monitoring.

However, even if we suppose that the relevant reasoning occurs at the level of discourse monitoring, the convention-based account would still only be plausible as an account of the (literal or nonliteral) *semantic* properties of utterances. The account would obviously still lack the generality we should expect. It is totally implausible that the beliefs and intentions regarding syntax and phonology required by a convention-based account of the syntactic and phonological properties of utterances should be present at the level of discourse monitoring or any other level. Therefore the Chomskian account should be our default hypothesis in each case, and we should abandon it only with great reluctance, given its theoretical virtues of simplicity and generality. This suggests that *even if the relevant Lewisian reasoning occurred* at, for example, the level of discourse monitoring, general considerations of simplicity and theoretical elegance dictate that we should not take this reasoning to *determine* the semantic properties of utterances, unless the Chomskian account is unavailable.

Furthermore, the level of analysis involved in computing nonliteral meaning or determining which utterance is conversationally appropriate (i.e., discourse monitoring) seems to require the prior analysis of literal meaning. This suggests that the relevant Lewisian reasoning plays no role in the processing of the literal meaning of utterances. The *non*literal meaning of an utterance is arrived at partly on the basis of the literal meaning of the utterance. And an appropriate response to an utterance is generated on the basis of beliefs about the environment and the person one is

[24] It is not clear to what extent these processes should be thought of as a single unified "level" of processing. Some aspects of our pragmatic processing may be relatively susceptible to such treatment, others not. Nothing I say turns on our being able to identify some particular "level" of processing, though I will continue to talk in terms of a "level" of discourse monitoring, for expository simplicity; I mean simply to refer to whatever aspects of our cognitive processing are involved in these processes.

[25] Though many of the difficulties raised in the last section would remain difficulties on these suggestions.

responding to (including her beliefs, intentions and expectations) *in conjunction with* the literal meaning of the utterance in question. Something like the literal meaning of an utterance is plausibly recovered from a representation of the syntactic structure of the utterance, together with recognition of the lexical items involved and retrieval of stored lexical representations. So the mental states posited by the convention-based theory do not seem to be involved in the processing of literal meaning. Thus if the relevant mental states do occur, they occur only after the literal semantic properties of an utterance have already been recovered.

Is it likely that the relevant Lewisian reasoning does in fact occur at the level of discourse monitoring? At this point some of the traditional problems for convention-based accounts that were briefly mentioned above seem relevant. Consider, for example, the case of a speaker addressing a hostile or unreceptive audience, or a cognitive agent producing utterances with no intended audience, perhaps talking to herself, or writing notes in a margin which she never intends to be read (see Chomsky 1975, Schiffer 1972, 1982, 1987). In both of these cases Lewis's account would seem to fail. Others do not respond to her utterances by coming to believe that they are true in *L*, and she does not believe that others are so coming to believe. Thus at least the first two clauses of Lewis's account are not satisfied.[26] To take a different sort of case, I might utter a sentence I believe to be false in metaphorical, sarcastic or hyperbolic uses of language. For example believing that Jones has done a really bad job, I might say in a sarcastic tone of voice:

(11) Nice job Jones!

Similarly, I might utter a sentence which I know to be false, as an example of such a sentence, or just to be silly, or as part of a story, or in an effort to deceive my audience, or just because I feel like it. I might produce (12), for example, knowing full well that it is not true:

(12) Ronald Reagan was born with no nose.

Despite the fact that Lewis's account is not satisfied with respect to my utterances in such cases (I don't believe them to be true for starters), they certainly have semantic properties; the utterances are not meaningless.[27]

[26] It is not clear to me which, if any, of the others are satisfied in these cases either. Clause 3, for example, seems to require belief that clause 2 holds. Similar points would seem to hold for the remainder of the clauses of Lewis's account.

[27] Lewis suggests three possible responses to the sorts of problems raised here (1983, pp. 183–4). The first is that we might simply stipulate that such "non-serious" language is automatically true in *L*. The second is that we might treat all such uses as exceptions. And the third is that we might define a certain set of core communication situations ("serious" ones) for which the conditions of the account must obtain. And this is all that would be required for a population to speak a language (it would not matter what would happen in the noncore cases). I am uncom-

I think it is clear that the Chomskian account being offered here most directly provides an account of the so-called "expression meaning" of an utterance. The basic idea is that an expression like (11), for example, is correlated in language processing with an LF representation, say, containing the mental words "nice", "job" and "Jones" in such a combination that the principles of compositional semantics governing these representations yield the truth condition, roughly, of a sincere utterance of (13):

(13) [You did a] nice job Jones.

The so-called "speaker/hearer meaning" is derived from this expression meaning.[28] The Chomskian has no need somehow to construct the expression meaning of an utterance out of possible speaker meanings. Rather the expression meaning is simply that meaning which is associated with the utterance in the course of language processing directly through principles of the compositional semantics embodied in the language processor. The speaker/hearer meaning is the intended/inferred ultimate meaning associated with the utterance, in this case through some further process of reasoning, presumably based on the utterance's expression meaning, and further facts. It seems that we might also think of this further meaning as being encapsulated in a mental representation with that content and inherited by the utterance in the Chomskian manner as well.

In the end, I think that it is unlikely that the conditions of any convention-based account are satisfied for any linguistic properties at any level at all. Even at the level of discourse monitoring, and even for semantic properties, there is substantial reason to doubt that the conditions of a convention-based account are satisfied. Certainly, speakers keep track of who their conversational partners are and attribute various speech acts to them. But this does not seem to require that the conditions of any convention-

---

fortable with the first sort of response because it further complicates the beliefs required by speakers to participate in Lewis's conventions. The other suggestions are possible moves Lewis could make. Of course we would still want the exceptions, or the noncore cases to *have semantic properties*. Presumably they must get them through their relation to core cases. The natural assumption would be that such utterances are relevantly of the same type as some core case. Perhaps they have the semantic properties that they would have if they were produced in some serious communication situation. There are at least two difficulties for this view. First, it may be that there are sound types corresponding to utterances which would *only* be produced in nonserious uses of language—sarcastic utterances of sentences, for example, often have different physical realisations than their nonsarcastic counterparts. And, second, we need some noncircular way of determining which situations are core/serious.

[28] I am not sure how much theoretical weight the terms "speaker meaning" and "expression meaning" can bear. They seem to me to be used in a variety of senses, which may or may not cohere. I only use them here to make contact with the literature.

based account of meaning be satisfied. It seems more plausible that the language processor computes various linguistic properties of utterances, while speakers are independently tracked at the level of discourse monitoring and assigned speech acts and other pragmatic aspects of the utterance are computed.

These considerations reinforce the suggestion that the Chomskian hypothesis should be the default hypothesis: it should be the default hypothesis in the case of the linguistic properties computed at the level of discourse analysis no less than in the case of linguistic properties computed at other levels of processing. Presumably the "output" of the discourse monitoring level of processing will be a representation of some (not necessarily literal) "meaning" of the utterance. Our default hypothesis should be that the utterance inherits this "meaning" in the Chomskian manner as well, regardless of what propositional attitudes play a role in the reasoning at the level of discourse monitoring which assigns such a representation to the utterance.

Considerations from reasoning at the level of discourse monitoring do not seem to provide any real help to the convention-based theorist. Such considerations offer no reason to believe that the empirical basis of the convention-based account is satisfied, for any linguistic properties at all. And the general theoretical considerations which establish the Chomskian account as the default hypothesis remain in place.

5. Conclusion

Where does all this leave us? If the language processor is thought of along standard lines as a special purpose cognitive mechanism, we can see language processing as basically an automatic process. The language processor can be seen as treating utterances as if they had their linguistic properties *intrinsically* in the sense that the processor does not take the arbitrariness of linguistic properties into account in processing utterances. In light of this, the various mental states which the convention-based theorist takes to underlie the semantic properties of natural language utterances seem strangely unmotivated.

It might be objected that convention-based theorists do not really intend to give a theory of processing at all. I think that this is right. My point however, is that if we do not have reason to believe that the basis of their account is satisfied by considering the states required for processing language, what reason *do* we have for believing that this basis is satisfied? The question becomes especially pressing when we see that an alternative account is available, and that there is good reason to believe that *its* empir-

ical basis *is* satisfied. If there is an alternative account of the nature of the semantic properties of natural language utterances, as I claim there is, the need for an account of these properties no longer provides us with good reason to believe that the empirical basis of the convention-based account is satisfied. And given that there is no independent motivation (from processing considerations) for believing that the empirical basis of the convention-based account is satisfied, why think that the convention-based account is true?[29]

Can independent motivation for the states underlying convention-based accounts be found elsewhere? Perhaps, for example, the convention-based account is more plausibly viewed as an account of how conventions get *established* rather than as an account of how they are *sustained.* Perhaps convention-based accounts might be required to explain language acquisition, or the evolution of language. More generally, there might be ways of making the Chomskian account and the convention-based account *compatible.* Perhaps the web of attitudes the convention-based account posits mediates the content relation for LFs. Perhaps it is necessary to account for the division of linguistic labour. Perhaps it is necessary to account for the establishing of sound meaning pairings (either developmentally or evolutionarily). If this were true, then we would not have to give up either account.

There are several points to make about these suggestions. First, and most importantly, I am perfectly willing to grant that the two accounts *may* be compatible. On the other hand, the bare possibility that the accounts can be made compatible shows very little. For all I know some theory incorporating phlogiston can be made compatible with current theories about oxidation. Maybe the phlogiston-based theory can be taken as a theory of some other vaguely related phenomena. On the face of it, though, we have two theories claiming to provide an account of the same thing—the nature of natural language semantic properties—and I claim to have argued for the conceptual and empirical superiority of one of these accounts. It hardly seems incumbent upon me to show that there is *no* way that the convention-based theory could be made compatible with this. Rather it would seem to be the burden of defenders of the convention-based account to find some interesting explanatory work for their theory to do. And in this I wish them luck. I claim only that it will not do as a

[29] Language processing and acquisition are without question central explananda of a theory of language. Arguably, we should look to such central explanatory roles of a given kind (asking what best explains them) in determining the nature of that kind. Since the Chomskian account ties the nature of linguistic properties directly to these central explanatory roles of linguistic kinds, it thereby seems to have a strong claim to being an account of the nature of these kinds.

theory of the nature of semantic properties of natural language utterances, since there is a better account of these.

Second, I am not claiming that propositional attitudes of various sorts concerning speakers and hearers or "the community" generally are irrelevant to semantics. It may be, for example, that the content of my mental representations is in some cases partly determined by the fact that I am willing to defer to experts. I see no reason why this is necessarily incompatible with the proposals of this paper, and so I can remain neutral on this question.

Third, and finally, although it is possible that the convention-based account can be made compatible with the Chomskian account, the defender of the convention-based account has to face the empirical and conceptual considerations raised above, and analogous considerations that might be raised for other possible explananda. Unfortunately I do not have space for a full consideration of these issues. However, I do not find any of the suggestions mentioned above at all plausible. The states underlying convention-based accounts seem to be as irrelevant to the process of language acquisition, for example, as they were to the account of language processing. The dominant view of language acquisition sees acquisition proceeding by means of the triggering of specific settings in an innate universal system of parameters. The acquisition of individual word meanings does not seem to require anything like the sorts of beliefs and intentions posited by the convention-based theorist either. Children's hypotheses concerning the possible meanings of words in their language are obviously highly constrained. While I am not really sure how the story would go for the convention-based theorist, the simplest formulations of such constraints do not seem to require the sorts of states that the convention-based theorist posits. (For some discussion of word learning see, for example, Markman 1989, Gleitman 1990, and Bloom 1993.) Regarding the evolution of language, very little, of course, is known. Though it is not really clear exactly what the convention-based theorist's account here is, an analogous position concerning, for example, the evolution of vision—that we must have entered into primitive vision conventions for vision to have evolved—just seems silly. And I do not see any reason to suppose that a convention-based account of the evolution of language should be taken any more seriously.

Much of the general theoretical and empirical support for the Chomskian view of language has come to light only fairly recently. So philosophers have not yet fully appreciated the implications of these results. In my view they provide powerful reason to believe that the convention-based approach is fundamentally mistaken and that the alternative Chomskian account is basically correct. Once we think in

terms of a special purpose cognitive mechanism for language processing, we see that it is not necessary to reason in terms of agents' beliefs and intentions regarding utterances and other agents' mental states in order to recover the semantic (or other linguistic) properties of utterances. It is not necessary to reason in this way, and there is no reason to suppose we do. Communicating in language is fundamentally different from signalling that I want to meet you over there by pointing vigorously in that direction.[30]

Philosophy Department
University of Manchester
Oxford Road
Manchester M13 9PL
UK
mfestsl@fs1.art.man.ac.uk

STEPHEN LAURENCE

REFERENCES

Bloom, Paul 1993: "Where Do Constraints on Word Meaning Come From?", in Clark 1993, pp. 23–34.

Chomsky, Noam 1965: *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.

——1975: *Reflections on Language*. New York: Pantheon Books/Random House.

Clark, Eve (ed.) 1993: *The Proceedings of the Twenty-Fourth Annual Child Language Research Forum*. Stanford, CA: CSLI.

Cooper, William and Walker, Edward (eds.) 1979: *Sentence Processing*. Hillsdale, NJ: L. Erlbaum Associates.

Curtiss, Susan 1988: "Abnormal Language Acquisition and the Modularity of Language", in Newmeyer 1988, pp. 96–116.

Fodor, Janet D. 1991: "Sentence Processing and the Mental Grammar", in Sells, Sheiber, and Wasow 1991, pp. 83–113.

Fodor, Jerry A. 1975: *The Language of Thought*. Cambridge, Mass.: Harvard University Press.

——1983: *The Modularity of Mind*. Cambridge, Mass.: MIT Press

——1990: "Psychosemantics: or Where Do Truth Conditions Come From?", in Lycan 1990, pp. 312–37.

Forster, Kenneth 1979: "Levels of Processing and the Structure of the Language Processor", in Cooper and Walker 1979, pp. 27–85.

Frazier, Lyn 1987: "Theories of Sentence Processing", in Garfield 1987, pp. 291–307.

[30] I would like to thank Eric Margolis, Stephen Schiffer, Gabriel Segal, Stephen Stich and the editor of this journal for helpful comments and suggestions.

——1988: "Grammar and Language Processing", in Newmeyer 1988, pp. 15–34.

Garfield, Jay, ed. 1987: *Modularity in Knowledge Representation and Natural-Language Understanding*. Cambridge, Mass.: MIT Press.

Garrett, Merrill 1990: "Sentence Processing", in Osherson and Lasnik 1990, pp. 133–75.

Gleitman, Lila 1990: "The Structural Sources of Word Meaning". *Language Acquisition*, 1, pp. 3–55.

Goldin-Meadow, Susan and Mylander, Carolyn 1990: "Beyond the Input Given: The Child's Role in the Acquisition of Language". *Language*, 66, 323–55.

Goldin-Meadow, Susan, Butcher, Cynthia, Mylander, Carolyn and Dodge, Mark 1994: "Nouns and Verbs in a Self-Styled Gesture System: What's in a Name?". *Cognitive Psychology*, 27, 259–319.

Grice, H.P. 1989: "Meaning", in Grice 1989, pp. 213–23. Originally published in 1957 in *The Philosophical Review*, 66, pp. 377–88.

Grice, H.P. 1989: *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.

Happe, Francesca 1994: *Autism*. London: UCL Press.

Jackendoff, Ray 1994: *Patterns in the Mind*. New York, NY: Harper Press.

Kasher, Asa (ed.) 1991: *The Chomskyan Turn*. Oxford: Basil Blackwell.

Lewis, David 1969: *Convention*. Cambridge, Mass.: Harvard University Press.

——1983: "Languages and Language", in Lewis 1983a, pp. 163–88.

——1983a: *Philosophical Papers*, vol. 1. Oxford: Oxford University Press.

Lycan, William (ed.) 1990: *Mind and Cognition*. Oxford: Basil Blackwell.

Markman, Ellen 1989: *Categorization and Naming in Children: Problems of Induction*. Cambridge, Mass.: MIT Press.

Marslen-Wilson, William and Tyler, Lorraine Komisarjevsky 1987: "Against Modularity", in Garfield 1987, pp. 37–62.

Matthews, Robert 1991: "Psychological Reality of Grammars", in Kasher 1991, pp. 182–99.

May, Robert 1985: *Logical Form: Its Structure and Implications*. Cambridge, Mass.: MIT Press.

Newmeyer, Frederick (ed.) 1988: *Linguistics: The Cambridge Survey*, vol 2. Cambridge: Cambridge University Press.

Osherson, Daniel and Lasnik, Howard (eds). 1990: *Language*. Cambridge, Mass.: MIT Press.

Petitto, Laura 1987: "On the Autonomy of Language and Gesture". *Cognition*, 27, pp. 1–52.

Pinker, Steven 1994: *The Language Instinct*. London: Penguin Books.

Poizner, Howard, Klima, Edward and Bellugi, Ursula 1987: *What the Hands Reveal About the Brain*. Cambridge, Mass.: MIT Press.

Prichett, Bradley 1992: *Grammatical Competence and Parsing Performance*. Chicago: University of Chicago Press.

Schiffer, Stephen 1972: *Meaning*. Oxford: Clarendon Press.

——1982: "Intention-Based Semantics". *The Notre Dame Journal of Formal Logic*, 23, pp. 119–56.

——1987: *Remnants of Meaning*. Cambridge, Mass.: MIT Press.

Sells, Peter, Sheiber, Stuart and Wasow, Thomas (eds). 1991: *Foundational Issues in Natural Language Processing*. Cambridge, Mass.: MIT Press.

Smith, Neil and Tsimpli, Ianthi-Maria, 1991: "Linguistic Modularity? A Case Study of a 'Savant' Linguist". *Lingua*, 84, pp. 15–51.

——1995: *The Mind of a Savant*. Oxford: Basil Blackwell.

Stabler, Edward 1983: "How Are Grammars Represented?". *Behavioral and Brain Sciences*, 6, pp. 391–421.

Swinney, D., Ford, M., Bresnan, J. and Frauenfelder, U. 1988: "Coreference Assignment During Sentence Processing". Unpublished Manuscript.

Wallman, Joel 1992: *Aping Language*. Cambridge: Cambridge University Press.

Zurif, Edgar 1990: "Language and the Brain", in Osherson and Lasnik 1990, pp. 177–98.