

## 10 Convention-based semantics and the development of language

---

*Stephen Laurence*

### 1 Introduction

A natural and historically popular view of the relation between meaning in public natural languages and thought is that thought is metaphysically prior, that natural language meaning depends on the meaning or content of mental states of one sort or another. Today this view is most closely associated with the work of Paul Grice and various philosophers broadly inspired by Grice's original work. This work (which I shall refer to collectively as 'Convention-Based Semantics') ultimately seeks to ground natural language meaning in a complex web of beliefs and intentions of language users. Many versions of Convention-Based Semantics have been put forward and, in spite of the numerous objections and purported counter-examples which have been raised against various versions of the theory, it remains the dominant view about the nature of natural language semantic properties amongst philosophers.

While I am sympathetic with these philosophers in their attempt to ground language in thought, I think that Convention-Based Semantics is wrong, not just in detail, but fundamentally. One basic way in which Convention-Based Semantics seems to me to go wrong is in assimilating natural language linguistic meaning to communicative phenomena generally. In my view it is a mistake to try to find a common reductive analysis that will ground natural language meaning and the meanings variously associated with pantomime gestures, with lighthouse beacon patterns, with hand gestures to help a friend back her car into a tight spot and any of a variety of other forms of meaningful communication. I think there exists substantial evidence for the special purpose nature of the language processor (see, for example, Bellugi *et al.*, 1993; Curtiss, 1988; Gopnik and Crago, 1991; Pinker, 1994; and Smith and Tsimpli, 1995), and the connection between natural language utterances and their semantic properties is likely to reflect the distinctive nature of this processor. Accordingly, the alternative general strategy I advocate attempts to ground natural language linguistic meaning in facts about principles and representations intimately

associated with the language processor (in the spirit of Noam Chomsky's accounts of the nature of natural language syntactic properties).

I believe that my alternative 'Chomskian' account is superior to Convention-Based accounts of natural language meaning in a number of ways. I think my alternative is simpler and independently well motivated, whereas I think the Convention-Based account faces substantial empirical and methodological difficulties (not just a few relatively technical counter-examples). Achieving a proper evaluation of Convention-Based Semantics is not an easy matter, however. There are a number of different versions of Convention-Based Semantics, and theorists disagree rather substantially over the proper interpretation of the theory and its guiding motivations. In an earlier paper (Laurence, 1996) I took the Convention-Based theorist to be largely concerned with how the connections between meanings and particular utterance types are sustained. A common response to my argument there has been that there needn't be any conflict between my Chomskian account and Convention-Based accounts; perhaps, for example, we should see the Convention-Based theorist as principally interested in how the connections between meanings and particular utterance types get *established*, rather than how these connections are *sustained*. This was an issue that I didn't have room to address in the earlier paper, and it is worth exploring since at least some Convention-Based theorists seem to be motivated by such considerations. Accordingly, in this chapter, my discussion will centre around the claim that considerations pertaining to the development of natural language (in language acquisition, or in the history of the species) provide some reason for adopting Convention-Based semantics.

The structure of my discussion will be as follows. First, in section 2, I will give a broad outline of the form Convention-Based accounts typically take, followed by some brief discussion of some philosophical motivations for pursuing the theory, in section 3. In section 4, I will argue that developmental considerations do not support Convention-Based accounts and that, in fact, quite the contrary is true: a variety of developmental considerations actually provide a relatively compelling argument *against* such accounts. In section 5, I present my alternative to Convention-Based Semantics, arguing that this account is fully consistent with the data discussed in section 4. In this chapter, as in the earlier paper I mentioned, some of the arguments I give are strongly based on empirical considerations. And I think many philosophers who advocate Convention-Based Semantics think that these kind of empirical arguments are largely beside the point for one reason or another. This seems to be the sort of view that Grice himself held (see Grice, 1976/80 and below). Though this sort of view is not uncommon, it is not exactly clear to me how or why philosophical theories of meaning or other phenomena should escape empirical constraint. In any

case, in the final section I discuss some philosophical strategies for evading such constraints which some might find appealing, and I argue for the minimal sort of empirical constraint on philosophical theories about meaning that they should at least be compatible with our best empirical theories of natural language and its processing, development and evolution.

Though I am mainly concerned in this chapter with the nature of natural language semantic properties, a related issue concerns the role of internal representations of natural language linguistic properties in cognition. Are such representations strictly limited to use in language processing, or do they play a wider role in our cognitive lives? In Peter Carruthers' (1996a) terms, this is the question of whether the communicative conception or the cognitive conception of language is correct. While Convention-Based Semantics seems committed to the communicative conception wherein natural language is basically limited to its role in communication (though see Devitt and Sterelny, 1987), my alternative position is compatible with either the cognitive or the communicative conception depending on the role in thought which the underlying representations posited by the account play. The arguments here may, however, undermine the motivations of those who accept the communicative view of natural language because they endorse the Convention-Based Semantics programme.

## **2 Speaker meaning, conventional meaning and Convention-Based Semantics**

Following Grice, Convention-Based theorists typically distinguish two sorts of meaning, often called 'speaker meaning' and 'conventional meaning'. Speaker meaning is characterised as *what a given speaker means by uttering a given utterance*. 'Speaker' and 'utterance' here are to be read very liberally, so that speakers include also writers and signers and indeed any agent producing any actions which might be said to have meaning. So, if I mean something by waving a flag at you in a distinctive manner, then I am thereby to be counted as a 'speaker'. 'Utterance' is to be interpreted similarly, so that my flag-waving action counts as an utterance. Speaker meaning attaches to individual acts of utterance, and so on different occasions two acts of the same type might very well have different speaker meanings. Conventional meaning, in the case of language, is the literal linguistic meaning which attaches to expression types.

Convention-Based theories typically seek to reduce conventional meaning to beliefs and intentions of speakers and hearers by first reducing conventional meaning to speaker meaning, and then reducing speaker meaning to patterns of beliefs and intentions. We can also think of this in terms of construction rather than reduction if we like. So Convention-

Based theorists could equally be thought of as trying to construct conventional meaning out of speaker meaning and speaker meaning out of beliefs and intentions of speakers and hearers.

The first stage in the Convention-Based Semantics construction of meaning involves the construction of speaker meanings out of particular sorts of beliefs and intentions of speakers and hearers. The following account of speaker meaning is representative of the sort of account given at this stage:

*Speaker meaning*

For any speaker, S, audience, A, utterance, x, and meaning, p, S's uttering x has speaker meaning p just in case S uttered x intending,

- (1) A to come to believe p
- (2) A to recognise that S intends (1)
- (3) A to fulfil (1) partly on the basis of his fulfilment of (2).

So, for example, my uttering 'Cats have whiskers' has the speaker meaning that *Cats have whiskers* just in case:

- (1) I uttered 'Cats have whiskers' intending for you to come to believe that *Cats have whiskers*, and
- (2) I intended for you to recognise that I intended for you to come to believe that *Cats have whiskers*, and also
- (3) I intended for you to come to believe that *Cats have whiskers* partly on the basis of recognising that I intended for you to come to believe that *Cats have whiskers*.

The link between speaker meaning and conventional meaning at the second stage in the project is often given in terms of the existence of conventions, where these are understood roughly in terms of David Lewis' account (Lewis, 1969, 1983). According to Lewis, a regularity, R (among a population) counts as a convention when the following conditions hold:

*Lewis' general account of conventions*

- (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, pp. 164–6.)

Some specific convention instantiating this general schema for conventions is then used to link conventional meaning to speaker meaning. The following example is representative of the sort of convention appealed to here.

*Conventional meaning*

A sentence, *x*, has the conventional meaning that *p* (among some population) just in case there is a convention to use *x* to speaker-mean *p* (see Davies 1996, p. 120).

This account of conventional meaning should be read with Lewis' general account of conventions in mind, so that we have, for example that everyone uses *x* to speaker-mean that *p*, and everyone believes that others use *x* to speaker-mean that *p*, and so on. Taking our earlier example, this means that everyone uses 'cats have whiskers' to speaker-mean that *cats have whiskers*, and everyone believes that others use 'cats have whiskers' to speaker-mean that *cats have whiskers*, and so on.

The full Convention-Based account is a rather complicated story. For example, if we spell out just the second clause of the convention here with reference to the account of speaker meaning given earlier, we find that, among other things, in order for 'cats have whiskers' to mean what it does we apparently need a fifth order attitude. You need to *believe* that I *intend* for you to *recognise* that I *intend* for you to *believe* that cats have whiskers. Of course the fact that the theory is complicated is in itself no objection to the theory. I don't see any reason why we should expect an account of the nature of meaning to be wholly uncomplicated when it is spelled out in full detail. But this range of empirical consequences will be relevant in determining whether developmental considerations argue for the theory.

### 3 The alleged priority of 'speaker meaning'

Michael Devitt and Kim Sterelny are among the philosophers who take developmental considerations to offer some support to Convention-Based Semantics. In their philosophy of language text *Language and Reality*, they argue that we should take conventional meanings to be built out of speaker meanings partly because we can have speaker meaning without conventional meaning but not vice versa. And they offer developmental considerations, among others, in support of this.

Devitt and Sterelny here seem to motivate their programme of analysing 'conventional meaning' in terms of 'speaker meaning' (and speaker meaning in terms of beliefs and intentions) by first noting that these two different sorts of meaning exist, and then asking which of the two we should take to be more primary. They write:

We shall suppose then that the distinction [between speaker meaning and conventional meaning] is real. Which sort of meaning is more basic or prior? (1987, p. 121)

They then argue that conventional meaning cannot be taken as more basic since there are cases where we have speaker meaning but no conventional meaning, which suggests to them that it is speaker meaning which is primary. Devitt and Sterelny cite a variety of different sorts of cases here, including:

- (1) Cases involving a slip of the tongue which produce nonsense strings like 'Can I morrow your dotes?' (for 'Can I borrow your notes?').<sup>1</sup>
- (2) Cases involving communication through gestures and mime in the absence of a common language.
- (3) Considerations involving 'the original development of language' (1987, p. 120)

Devitt and Sterelny's strategy here, though perhaps not uncommon, strikes me as rather strange. The assumption that if there are two types of meaning, speaker meaning and conventional meaning, one should reduce to the other seems unjustified to me. Consider the case of sense and reference, two types of meaning which play a role in a variety of semantical theories. Gottlob Frege, who introduced the distinction between sense and reference in his paper 'On Sense and Reference', used the example of the expressions 'the morning star' and 'the evening star'. Frege noted that there was a sense in which these expressions had the same meaning, in that they both referred to the same object (which, as it turns out, is the planet Venus). On the other hand, there is another sense in which they have different meanings, in that they have different 'cognitive contents' for us, and in particular, while it is trivial to say that 'the morning star is the morning star' it is not at all trivial to say that 'the morning star is the evening star'. The senses associated with these expressions correspond to the different ways in which we are conceiving of what turns out to be one and the same object (Venus) – the first presenting Venus as the last 'star' visible in morning sky, the second presenting Venus as the first 'star' visible in the night sky. Having distinguished these two types of meaning, though, we are not necessarily tempted to suppose that one of them (sense or reference) must be more basic and the

<sup>1</sup> Example from Fromkin, 1982, p.6.

other analysed in terms of it. And certainly the fact that there can be sense without reference (e.g., for the expression ‘the present King of France’) doesn’t provide us with a compelling reason to suppose that reference reduces to sense. We can, if we like, just recognise the existence of two different sorts of meaning, neither of which reduces to the other (though they are related to one another in various ways). Why shouldn’t we be able to say the same of ‘speaker meaning’ and ‘conventional meaning’?

Though the general principle that if there are two types of meaning and one can exist in the absence of the other, the second reduces to the first does not seem to be a valid principle, it might be that there is something about the particular case of speaker meaning and conventional meaning that makes this principle applicable. So we should look more closely at the cases of dissociation that Devitt and Sterelny cite. Unfortunately, their remarks are fairly brief, but about the development of language they say the following.

Consider the original development of language. Presumably this development was replete with examples of noises and gestures being used with communicative intent – with speaker meaning – before there existed a settled system of conventions for so using them. Communicative effort that was at least partly successful must have been a precondition for the development of linguistic conventions. The conventions came from regularities in speaker meanings. (1987, p. 120)

This passage is not easy to interpret, but Devitt and Sterelny seem to be claiming here that speaker meaning in some sense *had* to precede conventional meaning. We might try to reconstruct the argument here along something like the following lines. Conventional meaning is the shared meaning that expression types have amongst the members of a given community. But *shared* meanings could not exist until individual acts of meaning existed – indeed until a *number of individuals* meant the same thing by uttering tokens of a given expression type. Similarly, conventional meaning is by its very nature a meaning that is underwritten by the existence of a *practice* of meaning something by a given expression type. But *practices* require, by their very nature, more than one instance to become established. Much the same sort of argument might be made in the case of neologisms, and again perhaps in the case of language acquisition.

I don’t know if this is the sort of argument Devitt and Sterelny have in mind or not. But I don’t think it’s a good argument. The argument is certainly encouraged by the fact that Devitt and Sterelny use the terms ‘standard’, ‘literal’ and ‘conventional’ interchangeably in connection with this type of meaning. This conflation, however, begs some important questions, not the least of which is whether there is some interesting sense in which literal linguistic meaning is *conventional*. Certainly there could be no such a thing as shared meaning unless there were several individuals who meant

the same thing by their respective utterances of a given expression type. But shared meaning could as easily arise from the coincidence in *literal idiolect meaning* amongst several individuals as from speaker meanings. Similarly, though a practice of using a given expression may require more than one instance to become established, the instances on which the practice is based could well be instances of the literal use of a given expression type. So, since original meaning might just as well have been literal idiolect meaning, the argument does not show that speaker meaning must precede literal meaning.

Devitt and Sterelny's general argumentative strategy seems flawed in several ways, then. The existence of two types of meaning doesn't require one to be reduced to the other, so even if we could show that speaker meaning could exist without conventional meaning, but that the reverse was not true, this would not show that conventional meaning reduced to speaker meaning. And, moreover, it's not clear that the alleged asymmetry between speaker meaning and conventional meaning in fact exists.

#### **4 Do developmental considerations argue for Convention-Based Semantics?**

I want now to look more directly at some developmental considerations to consider their bearing on the plausibility of Convention-Based Semantics. Obviously I cannot undertake anything like a thorough review of language development as it might bear on Convention-Based Semantics. What I will do is briefly consider some relevant aspects of development and show that these aspects of development lend no obvious support to Convention-Based Semantics, and that some general features of the developmental problems in fact suggest the opposite conclusion.

##### *4.1 The development of language in the individual*

One place to look for developmental support for the Convention-Based account would be in language acquisition, and more particularly, in lexical acquisition. Perhaps the process of lexical acquisition crucially involves the patterns of mental states and processes the Convention-Based theorist posits (or provides more indirect support for these posits). Though the issue is extremely complex, not least because we are only just beginning to understand the processes involved in lexical acquisition, I will suggest two general sorts of considerations that argue against this hypothesis.

The first point is that the sorts of lexical acquisition principles currently under discussion in psychology do not seem to require any appeal to meta-beliefs. Psychologists studying lexical acquisition have posited the existence

of a number of general principles children employ in determining the meanings of new words. For example, Ellen Markman (1989) has proposed that children hypothesise word meanings which group objects taxonomically, as opposed to thematically, despite the fact that young children show a preference for thematic groupings of objects in non-linguistic sorting tasks. Lila Gleitman (1990) has proposed a rather different sort of constraint. She posits the existence of a syntax–semantics mapping, and suggests that children’s hypotheses about word meanings are constrained to respect the range of sub-categorisation frames associated with a given word. The syntax–semantics mapping determines a class of possible meanings for a word given the set of sub-categorisation frames associated with the word. Barbara Landau and her colleagues (Landau *et al.* 1988) have proposed that children hypothesise word meanings for count nouns that group objects according to shape. Applying principles like these, however, doesn’t seem to require any appeal to meta-beliefs at all. It is not obvious that the application of such principles even requires any awareness that other languages are spoken, or even possible. For all these principles seem to care, children could take themselves to be discovering non-conventional facts about a wholly invariant system used for communication. So the process of lexical acquisition, as governed by these principles, looks unlikely to support the Convention-Based account. Of course these principles don’t provide a full account of lexical acquisition, but they do underwrite a central aspect of meaning assignment where one might naturally expect the machinery in the Convention-Based account to show up if it was in fact essential to lexical acquisition.

The second point concerns the apparent dissociability of the acquirability of linguistic meaning from so-called ‘theory of mind’ abilities (our general capacity for reasoning about the mental states of others in terms of, for example, intentions, beliefs and desires). Evidence from recent studies of individuals with autism suggests that a central component of autism is the lack of a theory of mind (for a review of some of this literature, see Happé 1994). And, although most people with autism have poor or non-existent 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 quite significant linguistic abilities. Of particular relevance here, however, is the fact that people with autism seem to assign normal linguistic meanings to a range of lexical items. Helen Tager-Flusberg compared the extensions assigned to a class of concrete nouns by subjects with autism and those of normal subjects and mentally retarded subjects (all of the same mental age). She found no significant difference in results among the subjects in these different groups.

Autistic children at the same general level of vocabulary as control children do not acquire idiosyncratic word meanings. Rather, they show the same patterns of generalisation of meaning as evidenced by their overextension and under-extension errors. And indeed the patterns reflect adult judgements of these stimuli too. (1985, p. 1175)

Subjects with autism do differ from normal subjects regarding some lexical items – for example, words referring to mental states (as one would expect given their apparent deficit in understanding the nature of the states these terms refer to). But it is significant that they seem to be capable of assigning normal linguistic meanings to a relatively large class of terms (just how far Tager-Flusberg's results generalise is not clear). It is extremely puzzling, though, how someone lacking a theory of mind could assign normal linguistic meanings to terms (or use words meaningfully at all) if the Convention-Based theorist is right about the nature of meaning, given the numerous attitudes concerning the propositional attitudes of others that the theory posits. As we saw above, even for the simplest cases of literal meaning we would need as many as five orders of attitudes!<sup>2</sup>

Non-literal meaning may well require higher-order attitudes, as many pragmatics have suggested. And interestingly, subjects with autism do selectively have difficulties with such meanings. Indeed people with autism may provide an important source of evidence for testing theories in pragmatics, given their apparent theory of mind deficits. Francesca Happé (1993) has recently exploited just this property to provide an extremely interesting source of confirming evidence for Dan Sperber and Dierdre Wilson's (1986/95) Relevance Theory. The same sort of reasoning, however, suggests that lexical acquisition does not involve the patterns of attitudes which Convention-Based Semantics is committed to for ordinary literal linguistic meanings. (For some further discussion *see also* Hobson, 1993; Baron-Cohen, 1988).

In principle, people with autism could be considered exceptions to the generalisations in the Convention-Based account. However, this move strikes me as somewhat ad hoc. Certainly we should expect there to be exceptions to the generalisations here, but presumably they should be local and otherwise explicable – like slips of the tongue, for example. They shouldn't involve entire classes of individuals who clearly and systematically fail to satisfy the account for all uses of language. Indeed, it is not clear why, on a Convention-Based account, such individuals should be considered to be members of the linguistic community at all.

<sup>2</sup> Much the same point could be made using young children, since children do not employ propositional attitude concepts in an adult-like manner until the age of four (see Carruthers, 1996a, pp. 78–9). But presumably we do not want to deny that children before the age of four mean things by their utterances!

#### 4.2 *The development of language in the species*

Let's turn now briefly to the historical development of language in the species. Devitt and Sterelny describe a process of very slow development of language 'in humanoid society' by a process of 'lifting ourselves up by our own semantic bootstraps'. They suggest that at some point, early humans began to make noises with speaker meanings. When these 'caught on', conventional meanings were born. These became part of the culture, and were much easier to learn than they were originally to create. So energy could be focused on creation of further speaker meanings, which later would be conventionalised, and so on. For Devitt and Sterelny, it seems, natural languages are more or less the product of our general intellectual capacities applied to the linguistic domain. Their account of the development of language is not really an evolutionary account; language is a cultural artefact which required no language specific evolutionary adaptation to be produced.

Contrary to what this picture suggests, however, there seems to be good evidence that language *isn't* just a cultural artefact or human 'invention'. For example, there is no known correlation between the existence or complexity of language with cultural development, though we would expect there would be if language were a cultural artefact (Pinker, 1995). Further, language and general intelligence are dissociable. There are cases of normally intelligent people with extremely impoverished linguistic abilities, and there are cases of people with normal, even extraordinary, linguistic abilities despite severe general intellectual handicaps. The former sort of case is illustrated by the cases of Genie and Chelsea, neither of whom were exposed to a natural language until after the critical period, and whose linguistic abilities are severely impoverished (see Curtiss, 1988; Pinker, 1994). The latter sort of case is illustrated by the case of Christopher, studied by Smith and Tsimpli (1995). Christopher, despite having a non-verbal IQ of between 60 and 70 has normal English linguistic abilities, and moreover, has very impressive abilities in fifteen or sixteen other languages. Finally, languages that *are* the product of human invention are rejected in significant part by children learning language. Children exposed to pidgins – makeshift amalgams of several natural languages that are used for communication among a group of speakers with no single dominant language – reject them in favour of new languages of their own creation that are far richer than the pidgins they are exposed to (see Pinker, 1994). Similarly, deaf children studied by Susan Goldin-Meadow and her colleagues who were not exposed to a natural sign language, rejected the artificial manual communication system of their parents, in large part, substantially enriching and systematising it. So it looks like language is not a cultural artefact.

But if the Convention-Based theorist views language as a product of evolution and nevertheless looks to the evolutionary development of language for support for the Convention-Based account, then presumably the Convention-Based account would be about the development of language through precursors to language. It is certainly possible that some of the elements employed in the Convention-Based account (such as theory of mind abilities) formed a part of a precursor to language. But this is not the same as saying that the Convention-Based account itself was satisfied by some precursor to language. And, moreover, a precursor to language is not language itself – it is just a precursor (see Gómez, this volume). And we wouldn't say that because proto-vision (or proto-visual content) had such and such properties, vision (or visual content) must *reduce* to these properties. So it is not clear why we should say this in the case of language, even if the Convention-Based account were satisfied by some precursor to language.

## 5 A Chomskian account of natural language linguistic meaning

If Convention-Based Semantics isn't the right account of natural language linguistic meaning, what alternative account could be? The account that I favour is based on Noam Chomsky's general views on the nature of linguistic theory.<sup>3</sup> Chomsky claims that linguistics studies our knowledge of natural language, this knowledge forming a central and essential part of our capacities to acquire and process natural language. As with Convention-Based Semantics, there will naturally be a variety of different accounts in keeping with this general perspective on the nature of language. Focusing on linguistic meaning, we might provide something like the following sort of account, as a first pass.

### *Chomskian account of natural language linguistic meaning*

A sentence in my idiolect means what it does because it is assigned that meaning by the grammar that I have internalised.

Filling this out a bit, we might say that the linguistic meaning of a natural language utterance is given by the meanings associated (in the lexicon) with the words it is composed of (where these words are typed non-semantically – and presumably in virtue of analogous Chomskian account of phonological, morphological and syntactic properties<sup>4</sup>) and combined according to

<sup>3</sup> I call the account 'Chomskian' because I take it to be broadly within the spirit of Chomsky's views on linguistics, though Chomsky himself may well not endorse such a view. (For further discussion, see Laurence, 1996. However, I take the account here to be preferable to the more performance based account used there for expository purposes).

<sup>4</sup> It is worth noting here that, though these other features of utterances are equally conventional (in the pre-theoretic sense), no one believes that people have the requisite beliefs and intentions a Convention-Based account of these properties would require. One of the virtues of the Chomskian account proposed here is that it provides parallel accounts of all these linguistic properties. (For elaboration of this point see Laurence, 1996).

the syntax of the sentence and the principles of compositional semantics embodied in the system. So 'Cats have whiskers' means that *Cats have whiskers* in my idiolect because the grammar I have internalised assigns that meaning to this sentence. My internalised grammar assigns this sentence that meaning because it pairs a particular logical form with the phonetic form corresponding to 'Cats have whiskers', and the principles of compositional semantics governing my idiolect together with the semantic properties my lexicon assigns to lexical items associated with the logical form ('cats' 'have' and 'whiskers') yield the meaning *Cats have whiskers*.

This account is perfectly compatible with the various developmental considerations we've just been looking at. The sorts of principles governing lexical acquisition posited by Markman and others do not conflict with the account in any way. Indeed the process of acquisition such principles feed into just is a process which sets up the correlations between public language utterances and mental representations, in the lexicon and mentally represented grammar, which the Chomskian account appeals to. The account has no difficulties with the fact that some people with autism can acquire language and use it meaningfully, since the account does not appeal to sophisticated knowledge concerning the propositional attitudes of language users which people with autism seem to lack. And the various facts pertaining to the historical development of language in the species are perfectly compatible with the account as well. In fact, many of the considerations that were raised there are points among those typically cited in support of the general Chomskian model that I am appealing to. Finally, it is also worth noting that the machinery appealed to in the account is independently well motivated. We have strong independent reasons to posit the existence of a lexicon and to suppose that the lexical items associated with logical forms are governed by a compositional semantics.

This account has a number of other interesting features that I take to be advantages of the theory, which I would like to at least briefly mention. One advantage is that this account is idiolect-based. This allows for the fine grainedness of differences amongst idiolects, while still accommodating public language generality (as some function of idiolects). Since the account makes available a notion of literal meaning that is independent of public natural language meaning, such meanings can be constructed from the prior idiolect-based linguistic meaning (as other public language linguistic properties can be constructed from prior idiolect-based linguistic properties generally). At the same time, the community can enter the account of idiolect meaning via the meanings of the mental representations which fix idiolect meaning, if this is desirable (via deference to experts, for example). There are many interesting questions that arise in connection with these issues, which I do not have space here to pursue – about the relations between public languages and idiolects, for example, or the 'norma-

tivity' of meaning and language. These issues will have to await a future paper, however.

## 6 Deflationary interpretations of Convention-Based Semantics

Though Convention-Based theorists do often suggest arguments based on developmental considerations in favour of their accounts, they typically do not present the developmental considerations as offering direct empirical support for their views. Some Convention-Based theorists even seem to hold that empirical considerations, such as those pertaining to the development of language (in the species or the individual) are wholly irrelevant to the evaluation to their accounts. I find this sort of claim somewhat puzzling. I suppose there might be some construal of 'conceptual analysis' under which the theorist is simply interested in analysing our pretheoretic conception of 'language' or 'meaning' and so needn't look any further than our pretheoretic intuitions about how these concepts should be applied in constructing theories. But this seems to be a deeply uninteresting exercise. After all, for the most part conceptual analysis is interesting to the extent that it illuminates not merely our concepts but also the phenomena they pertain to. So, at the very least, the conceptual analyst's theories mustn't conflict with well-established empirical results.

There may be ways of interpreting Convention-Based theories that substantially reduce their empirical commitments, though. And I want to end by briefly discussing how Convention-Based theorists might try to do this. A number of theorists suggest, for example, that they are engaged in a project of 'rational reconstruction', and that such projects do not have the empirical commitments of, for example, reductive accounts of kinds in special sciences like biology. They do not often provide any clear account of what rational reconstruction involves, though, or why one would want to give a rational reconstruction, or when one is licensed. Dorit Bar-On's brief discussion of these issues in a recent paper (Bar-On, 1995) is perhaps the fullest discussion available. Bar-On argues for a 'genetic' interpretation of Convention-Based Semantics. Regarding the apparent empirical consequences of her view she says:

Now, of course, we are not to take this story as attempting to uncover actual historical conditions. It is not, for example, advanced as an empirical hypothesis of evolutionary biology. And the story would not be vitiated if it were somehow discovered that the actual facts did not bear it out. So what precisely is its status? My suggestion is that it should be seen as *a rational reconstruction of the condition under which language could emerge*. (1995, p. 97)

Bar-On distinguishes two kinds of 'rational reconstruction'. The first sort of rational reconstruction we might call 'Quasi-Empirical Rational

Reconstruction'. It attempts to provide a 'plausible' account of how a phenomena might come to be. Bar-On says of the Convention-Based account that while it doesn't purport to

track down actual historical conditions, [it] may still be taken to have quasi-empirical ambitions. By this I mean that it is possible (and even natural) to read the genetic story as a reconstruction of a path languageless creatures (like our distant ancestors) might plausibly take to get to language.

As Bar-On makes clear in a footnote, however, the accuracy of the account to any real process is not essential. Citing some potential counterevidence, she suggests that,

even in the face of such evidence, there is still room for a *plausible but false* story of the kind we have been telling here. Such a story might even be in some ways *better* – e.g., philosophically more interesting, illuminating, etc. – than a true empirical account according to which linguistic meaning emerged through evolutionary accident. (1995, p. 114)

Unfortunately, this is all that Bar-On says in clarification of this first variety of rational reconstruction. And it is clearly not enough. It isn't at all clear why a plausible but false story should be more illuminating or philosophically interesting than a true one (plausible or not). If the philosopher of language is trying to provide an account of linguistic meaning, and linguistic meaning is a natural kind studied in the special science of linguistics, it would be natural to suppose that accounts of the nature of this kind should be held to the same sorts of standards as accounts of the nature of other special science kinds. And presumably in providing accounts of the nature of biological, or geological, or economic kinds, plausible but false accounts are not worth much. At the very least, if linguistic theories and philosophical theories are both talking about the same subject matter – language – then philosophical theories should not conflict with well-established facts in linguistics (and related disciplines), just as philosophical theories in the philosophy of physics or biology should not conflict with well-established facts in those disciplines. The Quasi-Empirical variety of rational reconstruction doesn't seem particularly promising.

The other sort of rational reconstruction we might call 'Justificatory Myth Rational Reconstruction'. The idea here is to read the Convention-Based account as describing a mythical scenario for the development of language and to understand the role of the myth here as analogous to the role of the social contract myth in justifying political norms. The social contract myth describes a hypothetical state of nature and reconstructs the emergence of political norms in terms of people forming a social contract in the state of nature. For example, in John Rawls's extremely influential discussion, the 'original position' functions as the state of nature where people

make decisions about society from behind a ‘veil of ignorance’ (that is, without knowing what their place in the society will be). Here we can see the claim that some *actual* society is just (or not) as the claim that it is (or isn’t) the same sort of society as would be chosen in the original position. We are not committed to the historical accuracy of this process.

This, then, looks like it might be a promising way for the Convention-Based Semanticist to go. Unfortunately, it doesn’t succeed. To see why we need to look a bit more carefully at how the account is supposed to work for the social contract case. Here, an actual society is counted as just (or not) depending on whether a society of that sort would be chosen in the original position. How, though, do we determine whether two societies are relevantly of the same sort? If we say that they are alike in justice, this seems to presuppose that facts about justice are settled independently of this procedure. So presumably we choose other sorts of properties, properties such as distribution of material wealth, access to food, shelter and so on. When societies share properties of these sorts, then they are also alike in terms of justice, and so if a real society is like a mythical one in these respects, and the mythical one would be chosen in the original position, then the real one is just as well.

Turning now to the case of semantics, the problem is that there are no properties which can play the role that the properties of distribution of material wealth, access to food, shelter and so on play in the social contract case. The properties in question can’t just be non-semantic properties of utterances like shape or sound, for example, because something can share these properties with meaningful utterances and yet mean something completely different or nothing at all. Indeed this seems to follow directly from the very conventionality of language – the fact that our utterances in other circumstances might mean something entirely different or nothing at all. So, if a monolingual German produced a phonologically indistinguishable utterance to my utterance of ‘Susan leaped’ (‘Susan liebt’), their utterance would not thereby mean the same as mine. And, similarly, if my cat produced the sound ‘dog’, it presumably wouldn’t thereby mean *dog*. So sharing properties of sound or shape isn’t sufficient for meaning. And, of course, we cannot appeal to shared beliefs and intentions, because we are appealing to the Justificatory Myth precisely in order to avoid empirical disconfirmation. If we weren’t worried about all the relevant attitudes showing up in the actual account, there’d be no need to appeal to the Myth. So the Justificatory Myth version of Rational Reconstruction doesn’t seem to help the Convention-Based theorist either. It looks like rational reconstruction is not a promising strategy for reducing the empirical commitments of the Convention-Based account.

## 7 Conclusion

In this chapter we explored the possibility that developmental considerations pertaining to language acquisition in the individual or the historical development of language in the species provided support of some sort to Convention-Based accounts of natural language semantic properties. We found that in both cases, the developmental facts not only tended not to provide positive support for Convention-Based accounts, but actually provided evidence against such accounts. On the other hand, it was argued that the alternative Chomskian account advocated here is fully consistent with such facts. Faced with the threat these empirical results posed for the theory, Convention-Based theorists attempted to find a plausible deflationary interpretation of the theory. The ad hoc nature of the Quasi-Empirical version of rational reconstruction drove us to the Justificatory Myth version. But this last resort for the Convention-Based account was found not to work. The myth couldn't be connected up with the reality. Fortunately, a better alternative can be found. We can use the Chomskian account I've outlined above, and let natural language linguistic properties, including linguistic meaning, be inherited from properties of representations connected with the internalised grammar. And then we won't need the myth.

Thanks to Graham Bird, Peter Carruthers, Ted Elkington, Eric Margolis, Bernard Molyneux, Murali Ramachandran, Gabriel Segal and Dan Sperber for helpful comments or discussion of this material.