CHAPTER 36

MEANING AND REFERENCE: SOME CHOMSKIAN THEMES

ROBERT J. STAINTON

THIS chapter introduces three arguments that share a single conclusion: that a comprehensive science of language cannot (and should not try to) describe relations of semantic reference, i.e. word–world relations. Spelling this out, if there is to be a genuine science of linguistic meaning (yielding theoretical insight into underlying realities, aiming for integration with other natural sciences), then a theory of meaning cannot involve assigning external, real-world, objects to names, nor sets of external objects to predicates, nor truth values (or world-bound thoughts) to sentences. Most of the chapter tries to explain and defend this broad conclusion. The chapter also presents, in a very limited way, a positive alternative to external-referent semantics for expressions. This alternative has two parts: first, that the meanings of words and sentences are mental instructions, not external things; second, as Strawson (1950) stressed, that it is people who refer (and who express thoughts) by using words and sentences, and word/sentence meanings play but a partial role in allowing speakers to talk about the world.

Before turning to negative arguments and positive proposals, however, I want to clarify what the sources of the arguments are, and what view is being targeted.

36.0.1 The Source of the Arguments

The present chapter is not an attempt at Chomsky exegesis. Though many aspects of the arguments presented are explicitly due to Noam Chomsky, other parts of the argument(s) derive from the writings of other theorists, some thoroughly "Chomskian", some not: Akeel Bilgrami, Norbert Hornstein, Ray Jackendoff, James McGilvray, Julius Moravcsik, Paul Pietroski, and James Pustejovsky. For this reason alone, I describe the critical arguments and the positive alternative as exhibiting "Chomskian themes", rather than being exegeses of Chomsky's points. Of equal importance, the flow of the arguments—how the bits are put together into an overarching critique, and the formulation of the overall conclusion—are my own creations. Thus the arguments are at least twice removed from Noam Chomsky's writings.[1]

Having made the point that the topic is broadly Chomskian themes, let me say a bit more about what makes the themes "Chomskian". In brief, they all appeal, directly or indirectly, to *methodological naturalism*, "which holds that study of the mind is an inquiry into certain aspects of the natural world ... and that we should investigate these aspects of the world as we do any others, attempting to construct intelligible explanatory theories that provide insight and understanding of phenomena that are selected to advance the search into deeper principles" (Chomsky, 1993: 41). Three points about this orientation deserve emphasis. First, for the methodological naturalist, the object of study is a naturally occurring object, not an artificial construct. Second, she is seeking explanatory insight, not mere description (e.g. a systematic taxonomy is not enough). Third, the search for explanation is taken to imply, in turn, positing underlying realities, "deeper principles", that give rise to surface appearances; and to imply equally at least an attempt to connect up one's theories of language with other discoveries in the natural sciences. (Note: a genuine science of language, as understood here, involves the *aim* of *integration*; it need not involve *success* in *reduction*. These are quite different things.)

This methodological naturalist standpoint is taken to have several immediate and crucial implications. First, adopting methodological naturalism, the aim becomes not avowedly philosophical projects like "rational reconstruction", describing "knowledge which would suffice for interpreting", or defeating philosophical skepticism of some stripe, but rather the (familiar though terrifically difficult) project of empirically investigating, as best one can, the real (and frequently unobservable)

---

[1] Readers wishing to know what Chomsky himself thinks on these matters would do best simply to read him, especially the papers collected in Chomsky (2000a). McGilvray (1999) is also an excellent resource, which really does try to capture Chomsky's views and arguments, rather than merely presenting "Chomskian themes". See also Bezuidenhout, this volume. I should also make clear that, though the formulation of the conclusions is my own, as is the way the pieces of the arguments are put together, I am not presenting my own view on these matters.

features of a naturally occurring phenomenon, always looking forward to eventual integration with the core natural sciences.' Second, just as the sciences in general do not feel overly constrained by what "folk theorizing" suggests, or by how ordinary folks use words, for the Chomskian, linguistics and psychology are allowed to (indeed, they are expected to) put common sense conceptions aside, and to use terms in specialized ways, etc. Indeed, as Chomsky (1993: 25) rightly suggests, modern science gets going precisely when one is willing to be surprised by what are, from the perspective of common sense, "simple phenomena": e.g. that rocks fall, that people get sick and die, that a phrase is ambiguous, etc.[3] Third, since the sciences in general take their evidence wherever they can find it, there can be, for the methodological naturalist, no *a priori* restrictions on evidence in psychology or linguistics. On these grounds alone, much that has become conventional wisdom in the study of language — whether deriving from common sense talk, or from abstract philosophizing — has to be re-evaluated carefully.

If one studies the mind and language this way, taking preconceptions with a grain of salt, scientific inquiry into the salient natural object reveals — continues this line of thought — two less immediate implications. First, that the mature speaker/hearer's mind contains far more information than can be gleaned from the environment. This is the finding of the poverty of the stimulus. The most natural explanation of this finding, and the one that any unbiased scientist would immediately pursue, is that the human mind, including in particular the part of it responsible for language, benefits from a substantial innate endowment. A different though related hypothesis that emerges in this scientific endeavor is that the mind is divided, by nature, into a series of specialized faculties — rather than being, say, a homogenous "cognition/learning machine". This is the empirical hypothesis of modularity, with the language faculty being a case in point.[4] For the methodological naturalist, that some people find these latter results initially counterintuitive carries no real weight: after all, one should no more trust "intuitions" about brain structure and brain development than one should trust intuitions about the development and structure of the liver.

[2] See Chomsky, 1992a: 19, Chomsky, 1992b: 53 and Chomsky, 1994. For the idea of "knowledge that would suffice for interpretation", see Davidson, 1976, which builds on Foster 1976. For a very balanced comparison of this and other philosophical projects with Chomsky's naturalist one, see B. Smith, 1992. A trenchant critique of the former projects may be found in Antony 1997.

[3] This does not, of course, entail that Moore-style "common sense propositions" — e.g. "that there exist now both a sheet of paper and a human h a n d (Moore, 1939: 165) — should be rejected as false. As will emerge below, the methodological naturalist perspective does not *conflict* with common sense views about particular matters; rather, it pursues a different path entirely. As Chomsky (1995b: 138–9) writes: "It is not that ordinary discourse fails to talk about the world, or that the particulars it describes do not exist, or that the accounts are too imprecise. Rather, the categories used and principles invoked need not have even loose counterparts in naturalistic inquiry".

[4] It's worth noting that Chomsky employs a different notion of 'module' than, say, Fodor, (1983) does. Also, some read Chomsky as merely *stipulating* that linguistics, in his sense, studies what he labels "I-language": the intensionally characterized rules internal to the individual language faculty. This understates his claims. Chomsky's point, I take it, is that an unbiased methodological naturalist will study I-language, rather than other possible constructs, because the I-language construct turns out to correspond to a real aspect of the natural world that emerges in careful inquiry, whereas other constructs do not. Again, see Bezuidenhout, this volume, for more.

It is this standpoint, here called "Chomskian", that will be seen to call into question, in several different ways, the idea that a comprehensive scientific semantics should be in the business of pairing public language words and sentences with external objects, sets of external objects, and world-bound thoughts.

36.0.2   The Intended Target: Word–World Relations in Semantics

Having clarified what fundamental commitment underlies the rejecting, let me now clarify the view to be rejected.

It may be the default view in philosophy of language that natural languages are, at least in key semantic respects, rather like the formal languages invented by mathematical logicians. (That the logical languages are invented, with their properties being explicitly stipulated, is meant to be an unimportant difference.) This is the first plank of the view to be rejected.

In logical languages, like those invented by Frege, Russell, and Tarski, there are primitive formal elements of a few basic kinds (e.g. constants, predicates, quantifiers); and there are lexical semantic rules, which assign an interpretation to each primitive.[5] What this typically involves, at a minimum, is that (primitive) constants are assigned individual objects, and (primitive) predicates are assigned sets of objects. In addition, there are syntactic rules for recursively composing parts of the symbolism into larger wholes; and, there are corresponding compositional semantic rules which determine an interpretation for each resulting complex, given the interpretation of its parts and how those parts are combined. In particular, in some of these logical languages, whole sentences are compositionally assigned truth values as their meanings.

On the view to be rejected, natural languages are supposed to share all of the core interpretational properties of the logical languages: natural language names like 'George Bush', 'London', and 'Aristotle' are assigned real-world objects; natural language predicates (e.g. adjectives and verbs) are assigned sets of real-world objects (e.g. 'sings' is assigned the set of things which actually sing); and sentences are compositionally assigned... Well, here matters get rather complicated. Different philosophical theories notoriously make rather different claims about what should be compositionally assigned to natural language sentences. Even restricting ourselves to indicative sentences, it obviously won't do to assign natural language sentences truth *values* as their interpretation, as is done in some formal languages, for two obvious

---

[5] I use the word 'assign', in describing both the invented logical languages and natural language, without intending any specific account of what the relation "being assigned to" amounts to. As Ernie Lepore (p.c.) reminded me, some semantic theorists think of assigning as involving something like a function whose input is a formal item and whose output is an external entity. That, very roughly, is how Frege, Carnap, and Montague conceived of 'assigning'. But others, including Davidson and Tarski, take a more deflationary view of what I'm here calling "assignment". Since what matters in what follows is the status of the relata, not the nature of the relation between them, I gloss over these differences here.

reasons. First, not all true sentences mean the same thing, nor do all false sentences mean the same thing: 'Five is larger than two' is not synonymous with 'France is in Europe', though both are true. So truth values are not finely grained enough to be meanings. Second, it's not obvious that sentences are even the right kind of thing to be true/false. Certainly many sentences are explicitly context sensitive: 'He bought that yesterday', for instance, is at best only true relative to an assignment of values for 'he', 'that', and 'yesterday'. Thus, this sentence just isn't the sort of entity to even have a truth value tout court. To overcome the first problem, one can take sentences to denote not truth values but propositions made up of objects and properties; or one can take them to have a truth value as referent, but a Fregean Thought as sense; or one can take the meaning of a sentence to be truth-conditions. Such are the treatments of natural language sentences proposed by the contemporary Russellian, Fregean, and Tarskian respectively. These distinctions about sentence meanings, though important for other purposes, can be glossed over here: in the present chapter, I will speak vaguely of indicative sentences "expressing thoughts". To overcome the second problem, that natural language sentences contain context-sensitive items like 'he', 'that', and 'yesterday', one can say that such sentences express not thoughts exactly, but "proto-thoughts": something which is true or false relative to a set of contextual parameters (time, place, speaker, addressee). Such, then, is the first plank of the view to be rejected.[6]

Typically added to this idea, in "mainstream" philosophy of language, is the assumption that languages are the common property of a whole community, such that the symbolic items (words, predicates, sentences) are all public property. Languages, on this view, exist independently of speakers; and, being a public entity, each speaker typically has only a partial grasp of his/her shared language. This is the second plank of the view to be rejected. (For more on the Chomskyan target, see Bezuidenhout, this volume.)

36.1 THE THREE NEGATIVE ARGUMENTS

36.1.1 The Radical "Argument from Ontology"

Having clarified what the target is, and what makes the critiques in question broadly "Chomskian", I turn to the negative arguments.

The view at issue invokes relations between public linguistic items on the one hand (names, predicates, sentences) and worldly items on the other (external objects, sets, world-bound thoughts). Obviously, then, the relata must be able to stand in the

---

[6] Famously, the philosopher-logicians who are taken to be the grandfathers of this tradition—i.e. Frege, Russell, and Tarski — explicitly disavowed the idea that natural languages, in all their messy detail, could be treated this way. It was their contemporary philosophical followers, most notably Davidson (1967) and Montague (1974), who took the analogy between logical and natural languages literally. Chomsky et al. emphatically agree with the grandfathers, and disagree with their contemporary heirs.

requisite relations, including in particular (something like) the denotation relation. But, patently, the relata can't do this job if they aren't real. The first negative argument questions whether they are.

The widespread idea, to be questioned here, is that there are *public* signs, of shared languages, available to have referents: words, phrases and sentences that belong to languages like English, Urdu, and Swahili. In what follows, I will introduce three worries about this idea, to give the flavor of the thing. (One could easily raise many more.)

The first problem has to do with individuation of words given variation. Crucially, as actual working linguists are wont to note, the way we divide up languages in common sense, and in much philosophical theorizing, does not actually correspond to any robust divide. One speaks of "Chinese" as a single language, despite the fact that its two largest "dialects", Mandarin and Cantonese, are not mutually intelligible. In contrast, we call Spanish, Portuguese, and Italian different languages, rather than speaking of several dialects of Romance, just as we treat Swedish and Danish as different languages—this despite the fad that they are much more similar to each other, and far closer to mutual intelligibility, than the "dialects of Chinese" are. The only semi-robust divide here is mutual intelligibilityand, as noted, "languages" are not divided along those lines. One might reasonably reply that this worry can be overcome by thinking of words as belonging to dialects, not languages. But that won't really help, since what counts as a dialect is equally peculiar: Canadian English is supposed to be a single dialect, despite the many differences between speakers in urban centers and rural areas, and differences among the East, Central Canada, and the West; it also is supposed to be a different dialect than what is spoken in, say, Ohio. Clearly, we slice things as we do—both "languages" and "dialects"—not because of any robust linguistic divide, but because of colonial history, similar writing systems, shared canonical works of literature, present military might, arbitrary national boundaries, religious differences, and so on. That, and not "nature's joints", is what makes it the case that people "speak the same language/dialect". As Chomsky puts it, "This idea [of a common public language] is completely foreign to the empirical study of language... What are called "languages" or "dialects" in ordinary usage are complex amalgams determined by colors on maps, oceans, political institutions and so on, with obscure normative–teleologicalaspects" (1993: 18–19). (See also Chomsky, 1992b: 48; Chomsky, 1995b: 155 ff, and Bezuidenhout, this volume.) Instead of public languages/dialects, the real objects that one finds are (i) individual idiolects, (ii) sets of idiolects that share some non-obvious underlying parametric feature (e.g. having complements falling after heads), and (iii) the universally shared language faculty. None of these, however, corresponds even remotely to "public languages" like English and Urdu.

Now, to come to the problem of immediate interest here, if the boundaries around "languages" (or "dialects") don't reflect an objective difference in kind, what individuates *a word in a language?* What makes it the case, for instance, that distinct pronunciations are pronunciations "of the same word", if there aren't really objectively distinct languages? To take an example, why are 'fotoGRAFer' (said in Bombay)

and 'foTAHgrafer' (said in Toronto) the same word, yet 'fotOgrafo' (said in Buenos Aires) is not the same word as the former two? We are wont to say that there are *two* words here — the "English" word and the "Spanish" word — not three words. But this won't do, if "English" isn't objectively real: after all, all three differ in pronunciation. For that matter, even within a single country, or a single part of a country, there can be many "different pronunciations of the same word. So, as noted, appeal to local dialects isn't likely to help either. For instance, even within the Eastern United States, there are many pronunciations of 'Harvard'. More than that, children don't pronounce things the same way adults do, women don't pronounce things the same as men, and so on. Given variation, there thus seems to be no good reason to count public words the way common sense wishes to: we can't put aside the differences on reasonable grounds.

A natural reply to this first problem about counting words is that a dialect, or a language, is the symbol system shared by community such-and-such. But this reply is quite unhelpful, for at least two reasons. First, a specific worry: it's not possible to individuate the right community except by appeal to shared language. In particular, as we saw above, mutual intelligibility won't allow us to distinguish groups along lines that correspond to "languages". What "the community" for whom 'foto-GRAFer' and 'foTAHgrafer' are supposedly one word really have in common, and what distinguishes this "community" from others, is that everyone in it speaks English! A broader worry is that communities are no more "robust" than languages turned out to be. So even if one could divide languages in terms of which communities used them, this still wouldn't yield the kind of robust divide that the methodological naturalist demands.

There is a second reason why it is hard to individuate "public language words", beyond the problem of individuation in the face of across-speaker variation. It has to do with how to count words even granting the existence of languages/dialects. To pick an example essentially at random, is there one word 'forge' which has multiple meanings: create a fraudulent imitation, shape by heating in a fire and hammering, and *furnace* or hearth for melting or refining metal? Or are there three words, one for each meaning?' And, even restricting ourselves to one of the meanings, are 'forged', 'forges', and 'forging' wholly different words, or are they merely variations on the same word? What about the tensed verb 'forged', as in 'He forged the document', the past participle, as in 'He has forged many documents' and the adjective as in 'A forged document'? Are they precisely the same word, wholly different words, or variations on a single word? Also, if there is just one word here, or variations on it, what is that word? Rather than calling out for discovery of something real, these seem matters of decision.

In light of these questions about individuation, both across and within a "dialect", one can readily doubt that there is any such thing as "words in English", "sentences

---

[7] **And note the potential problem of circularity, if one does individuate words by their meanings. Meaning, recall, is supposed to derive from having two things stand in a relation. But now it turns out that one of the relata, on the "word" side, is individuated in terms of the other.**

in Swahili", and so forth. To echo Quine, one might insist that there can be no entity without identity. But if there are no such objects, there patently cannot be a science of word–world relations that pairs "public words and sentences" with worldly objects, sets, and proto-thoughts.[8] (Granted, for all that's been said so far, there might be *other* things that can be paired with external objects: morphemes of an individual's mental lexicon, for instance. But this possibility offers little solace to the kind of theorist that Chomskians are targeting.)

I said that I would introduce three problems about words. The third one involves issues about language norms. Though almost universally used among "English speaking" children, 'broked', 'runned', 'swirnrned' and so on are not "words in English"? Or again, despite its constant appearance in speech and writing, there isn't supposed to be a word in English that means it *is to be hoped that,* and is pronounced 'hopefully'. On the other hand, supposedly there is an English word pronounced 'ke-naw', because that's how Shakespeareans said 'know'; and there is, according to my Oxford dictionary, an English word 'peavey', even though almost no one would recognize it as such. These latter items aren't used, but they are "English words"; the former items are used, but aren't "English words". Clearly, what rules these words in or out is not how people *do* speak, but rather something about how they *should* speak. It's at least not obvious how there can really *be* such things, to stand in objective relations with external objects, sets thereof, and so on.

I pause to quickly summarize, before introducing a major objection to this line of argument. Because there is no objective way to individuate/count words (across or within a "dialect"), and because what makes something a shared, public word, if there really were any, would need to appeal to "ought" rather than "is", the Chomskian concludes that there aren't really any "public words". But then there cannot be a comprehensive science of language that pairs words (and sentences) with external things. Such is the radical argument from ontology.

A natural reaction to the claim that words (e.g. 'forged', 'photographer' and 'Harvard') are not real objects is perplexed disbelief that the claim has been seriously made. Surely it's just obvious that words exist. Besides, if an argument is needed, there is this: here we are discussing the various pronunciations *of the words 'Harvard' and 'photographer';* and above it was said that the word 'peavey' exists because of norms. But how can something which doesn't exist have different

---

[8] For those familiar with Chomsky's (1986) terminology, the central point may be put like this: public language words/sentences are part of the E-language picture, and the methodological naturalist must eschew E-language as not a suitable candidate for scientific study. See Bezuidenhout, this volume, for discussion.

[9] This point relates to another one that Chomsky regularly raises. Public language approaches are at a loss to find a "thing" which children under, say, six years of age know. They don't yet "know English" (or Swahili, or Urdu, or...). Indeed, there is no "public language" which they know at this age. But then how, positively, are we to describe the state of their minds? It seems absurd that we can only make the negative claim: i.e. that they do not yet know English (or Swahili, or Urdu, or...), but are on the way to doing so. Note too, how well such children communicate. This puts the lie to the idea that having "a shared public language" is genuinely necessary for communication.

pronunciations—as the argument itself grants that 'Harvard' and 'photographer' do? And surely, if something exists because of norms, then it exists. (Indeed, we seem to infer the non-existence of words on the grounds that 'peavey' exists! That's patently absurd.) Our discussion thus seems to give rise to paradox. Given the obviousness of the existence of words, and the paradoxes that quickly arise from denying their existence, it's hard to see how it could be suggested, at least with a straight face, that public words do not really exist.

There are several replies to this natural worry. On the one hand, one can agree that these things are real enough, but go on to question whether there could be a *science* that treated of them. Where by 'a science' is meant, to repeat, rather more than "any inquiry that is both theoretical and empirical". As hinted at the outset, 'science' in the context of methodological naturalism means, at a minimum, seeking explanatory insight; which in turn entails positing underlying realities, and aiming for integration with the core natural sciences. Many things exist which are not subject to scientific investigation, in this sense. This concessive reply will be considered at length in the second negative argument. To anticipate briefly here, the core idea is that the standard for being a "real object" has been set too high in the discussion above. It's not just public languages and words, but corporations, songs, countries, universities, national dishes, hair styles, TV shows, etc., that won't really exist given this over-high standard. Indeed, it's arguable on similar grounds that none of us exist: to see why, think of the enormous puzzles about how to individuate persons. A natural alternative view, which doesn't set the standard so high, is that perfectly real objects can be quite hard to individuate/count, and can be norm-bound. They need not require a "robust divide", but can rather be objectively different only in degree, with human interests setting the kind-divide between them. One could thus allow that there is such a thing as English (and other public tongues), and that the nature of English and the words/sentences in it depend on a host of complex relations (political, military, historical, religious, etc.)—including even explicitly normative ones having to do with "correct speech". Adding, goes the reply, that this does not make English and its elements unreal. Personally, I think there is something very importantly right about this. Still, the key point that will re-emerge in negative argument two is that, even granting this, one is hard-pressed to rescue the idea that a genuine *science* of language can, or should try to, describe word–world relations. Indeed, the account proposed of what makes words and languages real—e.g. that their individuation rests on norms, quirky anthropocentric interests, and a complex mess of other things—pretty much ensures that they will *not* be scientifically tractable. More than that, if that's what makes something a "word", it's not even plausible that "public word" will be an *idealization* that *will* be of any use in science. As Chomsky puts the general point:

Such informal notions as Swedish-vs.-Danish, norms and conventions, or misuse of language are generally unproblematic under conditions of normal usage, as is "near New York" or "looks like Mary". But they can hardly be expected to enter into attempts to reach theoretical understanding. (1993: 20)

As I say, this concessive reply will be elaborated at length in the next section. But there are non-concessive replies too, which try to defend the radical version of the "argument from ontology" according to which one side of the supposed relations (i.e. the public words/sentences) just do not exist at all. Let me introduce a couple of those replies here. That there are no public words or languages strikes us as absurd, but—goes the first reply—that is because we are taken in by an illusion of some sort. Part of the concept of "public word", the argument would go, is that the things in question are "out there", the shared property of many. They are not inside the mind. Given this, the public word 'Harvard', the story would go, is "unreal" in roughly the same way that the sky, the daily sunset, perceived colour, and rainbows are not real *considered* as *external objects.* In all these cases, we project "out there" something that is really an amalgam of things going on inside the mind, and (non-obvious) things that are going on in the external world: "the structure of language is not "out in the world but [is] rather a consequence of the mental organization of language users" (Jackendoff, 1987: 133). Ordinary people cannot fail to *think* of the sky, the sunset, blueness, and rainbows as mind-external objects, wholly out in the world, even after careful scientific training. But what scientific investigation teaches is that, appearances notwithstanding, they are partly in the individual mind. (Importantly, being open to taking these results seriously, thereby setting aside common sense, is part and parcel of being a methodological naturalist.)

The illusion that there really are public words, words "out there" that we share, is reinforced by the fact that people talk about words. An egregious case in point, as noted: the very argument against the existence of words apparently used as premises *claims about words.* But, coming to the second reply, that we talk about, say, 'Harvard', does not actually entail that there is a public word "out there" that we share. On the one hand, speakers regularly refer to things that simply do not exist: Santa, unicorns, the present King of France, etc.[10] On the other hand, even if there are some unquestionably real things that we refer to, when we speak of the word 'Harvard', there needn't be a single object which is *the* public word. *A* plausible alternative view is that there are many, many words 'Harvard'. For some purposes, we count all pronunciations as constituting "the word 'Harvard' "; for other purposes, we count only very few. And so on. We refer to different sets on different occasions, depending upon the context. The resulting sets are real, and they are intersubjective. Still, there isn't one thing, the word 'Harvard'. (See Bilgrami 2002 for this general line of thought.) Hence we can consistently talk about "the different pronunciations of 'Harvard' ", without committing ourselves to there being one unique thing, that publicly shared word, that can stand in a refers-to relation.

If the foregoing considerations work, then there cannot be a comprehensive scientific semantics that treats of relations between public words and external things. Indeed, what makes this the "radical variant" is that if this criticism is successful,

---

[10] Put metalinguistically, reports of speaker reference are referentially opaque in a way that expression reference, if it existed, would not be. See Bencivenga, 1983 for extended discussion; see also Jackendoff, 1987: 127 and Chomsky, 1995b: 150.

not only can there be no *comprehensive science* of word—world relations, there can be no truths *of any kind* that state relations between public words and worldly entities. (How could there be, if there aren't any public words?)

36.1.2 The Moderate "Argument from Ontology" (Science and Common Sense)

The first "argument from ontology" involved arguing that public words don't exist at all. The second "argument from ontology" accepts the reality of both relata. But it questions whether there can be a genuine *science* of the kind of common sense objects involved on both sides of the relation."

Crucial to the argument will be the contrast between the world revealed to us by common sense, and that revealed by modern scientific inquiry. We therefore need a way to draw that distinction. The methodological naturalist thinks there is a way to draw it, given nativism and modularity: we can distinguish the world revealed by common sense from the world revealed by science *cognitively,* in terms of the kinds of concepts deployed.

The concepts of common sense, in the sense intended here:

- are not social artifacts, but are rather part of our biological endowment;
- more precisely, they are constructed from innately given semantic features—though only the elements out of which the concepts are constructed are innate (and universal), not the resulting wholes; (See Chomsky 2000b: 185.)
- are acquired (rather than *learned),* and they do not need to be taught—indeed, given the poverty of stimulus, it's unlikely that they *could* be learned by/taught to a creature lacking the requisite innate endowment;
- are at the disposal of every non-pathological human;
- bring with them a rich and complex internal structure that eschews elegance in favor of day-to-day practicality, especially for living in human company—precisely because they are built out of an innately given store of features;
- have, finally, and related to this last point, inherently built in implicit references to human hierarchies, rights/obligations, and our intentional states, rather than aiming for an objective description of the world, independent of us.

Scientific concepts, in sharp contrast, *are* social artifacts.[12] More than that, a useful scientific concept is often a hard-won achievement of many years of collective labor. Such concepts must be taught; and frequently enough they cannot be learned, even by non-pathological people. Their content is austere, rather than rich. And, far from

[11] That the issue is a *science* of public words, and public languages, is missed by some of Chomsky's critics. See, for example, Wiggins, 1997.

[12] I am unsure whether Chomsky himself would endorse what follows. He sometimes suggests that humans have a "science forming faculty", and if scientific concepts derive from it then they are not especially social after all. Since Chomsky exegesis is not my aim, however, I leave this issue aside here. (Thanks are due to Julius Moravcsik for drawing my attention to the issue.)

being tied to anthropocentric interests, the whole idea of a scientific concept is to capture how things "really are" independent of us.

It is telling, too, that science involves explicit reflection not just on the concepts it creates to describe and explain, but also on what counts as good evidence, justification, etc. Those standards of evaluation too are sanctioned by groups, over extended periods of time; they aren't just "given by nature". Sciences, and scientific concepts, are thus artifacts of a social practice, rather than being innately specified — which allows, as Chomsky (1993: 32) suggests, that science can afford to disregard common sense, and is happy to move beyond some of its tenets.

The overall picture can be summed up with the following long passage from Chomsky's Language and Thought:

We have, by now, fairly substantial evidence that one of the components of the mind–brain is a language faculty, dedicated to language and its use—where by "language", now, we mean human language, not various metaphoric extensions of the term. Other components provide "common sense understanding" of the world and our place in it... Other components make it possible for humans to conduct scientific and mathematical inquiry, and sometimes to achieve remarkable insight: we may call them "the science-forming faculty", to dignify ignorance with a title. These could be quite different in character from those that yield "common sense understanding" in its various forms. It is an open empirical question, and no dogmatism is in order. The history of modem science perhaps suggests that the distinctions are not trivial; at least, that is one way to interpret the startling conflicts that have arisen between common sense understanding and what scientific inquiry reveals. (Chomsky 1993: 34–5)

Having contrasted science and common sense, we can now note that the concepts fall into two families, and add that each collectively provides a *perspective*.[13] What does the "scientific perspective" show us? Quarks, tectonic plates, genomes, and many other things. (To be clear, the scientific perspective does not merely reveal so-called "physical" objects; the mind–brain, at various levels of abstraction, can also be seen from this perspective. That, indeed, is just what linguistics is supposed to help reveal. And, of course, the common sense perspective does not merely reveal psychological states: it affords views of desks, house pets, tea, and toys.) Importantly, however, there is lots that the scientific perspective does not show us: it does not encompass normative categories like good wine, liveable cities, or well-prepared osso bucco; nor does it even encompass not-explicitly normative yet mind-dependent things like clouds,[14] tea,[15] desks, sunsets, breakfast cereal, and hockey scores. In

---

[13] Chomsky (1993:48) writes: "The information provided by lexical items and other expressions yields perspectives for thinking and speaking about the world..." Or again, "a lexical item provides us with a certain range of perspectives for viewing what we take to be things in the world, or what we conceive in other ways; these items are like filters or lenses providing ways of looking at things and thinking about the products of our minds" (Chomsky, 1992a: 36).

[14] On the natural assumption that whether a quantity of water in the atmosphere is a cloud depends upon mind-dependent relations — lie being visible, in normal circumstances, to the naked eye of normal humans.

[15] As Chomsky has frequently noted, what is chemically the very same substance could be tea — if created by dipping a bag of tea leaves into a cup of hot water — or contaminated water — if created, say,

general, being objective and ignoring interest-relative distinctions, the "scientific perspective" cannot see entities whose individuation conditions inherently involve complex human interests and purposes. Those things are only "seen" through the in-born lens of natural language expressions and the sorts of sub-lexical concepts that help make up these expressions.[16]

The reason why the scientific perspective cannot "see" such objects is most easily illustrated with examples. Take London. Chomsky (1993: 22–3) writes:

> We can regard London with or without regard to its population: from one point of view, it is the same city if its people desert it; from another, we can say that London came to have a harsher feel to it through the Thatcher years, a comment on how people act and live. Referring to London, we can be talking about a location, people who sometimes live there, the air above (but not too high), buildings, institutions, etc., in various combinations. A single occurrence of the term can serve all these functions simultaneously, as when I say that London is *so* unhappy, ugly, and polluted that it should be destroyed and rebuilt 100 miles away.

As Chomsky goes on to conclude, "No object in the world could have this collection of properties" (1993: 23). To be clear, it is not being claimed that London simply fails to exist. "London is not a fiction" (Chomsky, 1992a: 37). This is a key difference between the radical worry in the last section, about public words, and the present moderate worry. The issue this time is not whether the relata exist, but (as it were) what makes them exist: the worry is that the kind of socially constructed object that is London (and 'London'!), so highly dependent on human perspectives and interests as it is, cannot be seen by the peculiar instrument that is natural science. This, and not the very existence of London, is what seems doubtful. To generalize the point, imagine natural science looking for external world correlates of 'bargain', 'owner-ship', 'tenure', 'delicious', 'sacred', 'funny', 'notary public', 'nearby' or 'polite'. How can genuine science see any of these? The thing is, if many, or most, or even all common sense concepts are covertly like 'London' and the rest — and, given the contrast between scientific concepts and common sense concepts that emerged above, this is just what seems to be the case — then a scientific reference-based semantics is hope-less for the vast majority of lexical items in natural language.

Crucially for present purposes, granting that what common sense "sees" is per-fectly real, we still arrive at the conclusion introduced at the outset: that a com-prehensive science of language cannot (and should not try to) describe relations of semantic reference, i.e. word–world relations. That is because the things which man-age to be, on this more moderate view — i.e. both words and ordinary objects — are

by the right combination of organic matter falling into a lake. Because "what is tea" reflects human interests in this way, science isn't in the business of contrasting tea from non-tea. See Chomsky, 1995b: 128 and Chomsky, 2000b: 189 for discussion.

[16] Some of the authors discussed here — e.g. Jackendoff and McGilvray — take an additional neo-Kantian step, and draw the metaphysical conclusion that there are two "worlds", one for each perspective. But this is not an immediate consequence of the present argument about the nature of scientific semantics. One can resist the metaphysical step, and stick to the epistemological claim that the scientific perspective can only "see" part of the one world. Either approach will yield the desired conclusion.

nevertheless not real *in the right sort of way*. Hence they cannot be "seen" from the scientific perspective. They are as "real" as governments, townships, by-laws, nearness, corporations, national dishes, and so on — which, as Heidegger and *G*. E. Moore rightly insisted, are as real as can be — but, like townships and nearness, they are invisible to the working scientist. In particular, then, they are invisible to the scientist of language, as such.

Before continuing, I should consider some natural objections. As a reminder, I have been trying to contrast two ways in which to look at the world: from the perspective of common sense, and from the perspective of modern science. I did so by contrasting two families of concepts, namely common sense concepts, and scientific concepts. I then suggested that, from the scientific perspective, there are many things which just cannot be seen — and hence that no science of language can describe them. Consider now four objections to this line of argument.

First, one might object that the distinctions between the kinds of concepts are being drawn, at least in part, along lines that a radical empiricist or cognitive holist would question. The reply to this is obvious. As hinted right at the outset, methodological naturalists consider radical empiricism, and also anti-modularity, to be empirical theses. They also consider modularity and a very significant innate endowment for language to be well supported on empirical grounds. So, the fact that the case against reference-based semantics may rest on such hypotheses will not detain them.

Second, one might reasonably complain that there will be a host of potential referents which the natural sciences *can* see, and which therefore could be described in a theory of word–world relations (assuming public words to exist). These are precisely the objects that the perspective of science specifically affords: quarks, tectonic plates, wh-traces, etc. Thus the foregoing argument does not on its own rule out a "linguistic science of reference" *for the vocabulary of the sciences.* (See Chomsky, 1992a: 42 – 3 and Chomsky, 1993: 27 ff for discussion.) Granted, one might not be able to have a comprehensive science of language; but a limited reference-based scientific semantics would still be possible. However, this is ruled out independently, given the Chomskian view that the proper object of study, for a science of language, is the human language faculty. That, after all, is the aspect of reality that we find, when we start to "investigate language" naturalistically. Now, the language faculty is the innately given part of the mind–brain which, in response to environmental triggers, settles into a steady state of linguistic competence *sometime before age five;* crucially, then, not everything one "learns about language", as we pre-theoretically say, belongs in the language faculty. In particular, Chomsky suggests that scientific vocabulary likely is *not* stored in the language faculty: it is learned, not acquired; it is austere, not rich; it is a construct, not an aspect of our biology. Thus, while there might conceivably be a science of reference for scientific terms, it would not be part of the science of natural language.[17]

---

[17] **An interesting side effect of distinguishing common sense concepts from scientific ones is that one could follow Quine in holding that there is no analytic/synthetic distinction for *the* concepts used in**

Third, one might doubt that science and common sense really do provide mutually exclusive and exhaustive perspectives. First, it's clearly true that the supposed gulf between science and common sense *seems* bridgeable in places, and this on at least three grounds: (a) there seem to be factual claims that straddle the divide: e.g. 'My cup of herbal tea boiled at 101.35 degrees centigrade'; (b) common sense reports can clearly serve as evidence for/against a scientific hypothesis: e.g. 'The Prime Minister died after eating one of those' or 'It turned red like a fire truck when we poured the liquid on it' can both support (or call into question) a genuine scientific hypothesis about the chemical make up of an unknown item; (c) it's arguable that there are whole disciplines which straddle the borderline: criminology, epidemiology, anthropology, medicine, horticulture, etc. Thus the difference may be not a difference of kind, but one of degree. (See Moravcsik, 1998: 127 for extended discussion.) This is a very important objection. It therefore merits a detailed rebuttal, comprising three related replies. The first reply is that there cannot be a *comprehensive* science of language unless that discipline can see all (or at least the vast majority) of objects that can bear names. Hence the conclusion stands even if the difference between what common sense can see, and what modem science can see, is a matter of degree, with intermediate cases along the way—as long as there are lots of things on the extreme end that science can't see. In short, the conclusion argued for doesn't really require the stronger premise that the distinction is exclusive and exhaustive. The second reply makes a related point about making do with a weaker premise. What the conclusion requires is not really the claim that "science simply cannot see common sense objects"; what it requires, instead, is merely that there is no *single* science which can see (almost) every common sense object—since a comprehensive science of language that described word–world relations would need to be just such a science. To imagine a comprehensive science of language that posits word–world relations is, for instance, to envisage a *single* genuine explanatory science that can "see" all of friends, yarmulkes, Tuesdays, symphonies, jokes, vagrants, bargains, and every other common sense object.[18] Put another way, the second reply is this: what is required, at a minimum, for a comprehensive scientific semantics that introduces word–world relations is *There exists a science x such that, for almost every y, x can see y*; but at best what is plausible is *For almost every y, there exists a science x, such that x can see* y. The third reply to this third objection is that even this weaker claim (which wouldn't actually avoid the conclusion in any case) gets much of its plausibility from loose terminology. To repeat, as the term

science (since those words really do get their meaning holistically, from their place in a world-describing theory), while nevertheless insisting that common sense concepts—built as they are from innately specified features—will license analytic truths. See Chomsky, 2000b: 186. For a discussion some of the epistemological implications of this way of slicing things, see Matheson and Stainton, (2002).

[18] Jackendoff draws a weaker (but still very interesting) conclusion, on related grounds. He argues that "language is about entities in *the* world as *construed by the language user/perceiver*" (1987: 128). (See also Jackendoff, 1991: 12.) As a result, word–world relations cannot be studied prior to, and independent of, psychological investigations about how humans categorize. In which case, external-referent semantics cannot be a genuine alternative to Jackendoff-style Conceptual Semantics, since the former implicitly presupposes the latter.

is being used here, not every systematic empirical inquiry counts as 'science': e.g. an exhaustive taxonomy of the Earth's beetles is not science, in the sense I have been discussing — not least because mere taxonomy does not seek out underlying explanatory realities. Noting this, it's not even clear that most common sense objects are "seen" by any genuine science after all, though it might well be plausible that some systematic empirical inquiry can see each such object.

Whereas the three previous objections were specific to the moderate argument from ontology — each attacking in a different way the proposed sharp and exhaustive contrast between the perspective of common sense and the perspective of modern science — the final objection, which is methodological in character, applies more globally to both arguments from ontology. The complaint here is that even making these two ontology based objections flies in the face, paradoxically enough, of what I described as the core tenet of the Chomskian approach to language, viz. methodological naturalism. From that perspective, the test of a theory is, surely, whether it can establish a rich body of empirically supported doctrine; and, the fourth objection continues, semantics has made great progress in those terms. There are thick textbooks full of results, journals packed with data and detailed debates, and so on. Thus, the fact that reference-based semantics might fail to meet certain arch ontological scruples should not cause us to reject it, since it is a thriving research program.

Actually, Chomsky himself agrees with the methodological point: questions of ontology are, for him, posterior to questions about explanatory and descriptive success (Chomsky, 2000b: 184). Presumably every methodological naturalist will agree. The objection still misfires, however, because it is based on the (thoroughly mistaken) idea that "Chomskians" reject semantics root and branch. It's not the entire sub-discipline of semantics that is being rejected, only a peculiar spin on it. And, continues the reply, the many existing results of semantics have little or nothing to do with the extraneous philosophical hypotheses that shared public words/sentences stand for real-world external objects, sets thereof, and world-bound thoughts.[19] Indeed, meaning broadly construed remains as central to Chomskian linguistic theorizing in 2003 as it was in 1955.[20] To give but one example, Chomsky's most recent Minimalist theorizing makes essential use of the principle of Full Interpretation, which requires (among other things) that only elements that have an interpretation can remain at the end of a derivation — this being the point where

[19] Some would say that, beyond leaving the key results of formal semantics standing, absolutely *nothing* is lost by linguistics, when it abandons the reference-based approach; moreover, much is gained. Thus Hornstein (1989) writes: "If semantics is concerned with truth conditions, and this is construed as correspondence, then I can see no reason for thinking that there is *any* link between semantics so defined and theories of linguistic interpretive competence. Moreover, this is all for the good as far as the latter enterprise is concerned, for semantic theories seem to require the ascription of powers and capacities to native speakers which are as mysterious as those capacities that we wish to explain. Syntactic theories, those types of theories that eschew language–world relations, are not similarly problematic. It is for this reason that syntactic theories are *methodologically* preferable". See also scattered remarks in Hornstein 1984.

[20]   Neil Smith (1999: 163) quotes Chomsky as saying that "putting aside phonology, virtually everything I've done in the field since LSLT *[The Logical Structure of Linguistic Theory]* falls within semantics". See also Chomsky, 2000b: 174.

the string reaches the interface with Conceptual Structure. All other structure (e.g. nominative and accusative case marking, expletives, agreement features not bearing content) must be "checked", thereby being rendered "invisible", before this point. Thus meaning couldn't be more central to current linguistic theory. (See Radford, 1997: 170ff for introductory discussion, and Chomsky, 2001; Chomsky, 2000c; Chomsky, 1995a for the evolving details of Minimalism.) Moreover, as Pietroski (2003, 2005) points out, there is much of the same positive hard work for the sub-discipline of semantics to do, without reference and truth. First, this sub-discipline will explore semantic properties of expressions such as: what they can and cannot mean; whether they are ambiguous; if they are ambiguous, why, if they are not, why not; what referential dependencies must, can and cannot obtain; and so on. In fact, for all that has been said here, semantic theory could even retain the architecture of a Fregean or Tarskian theory, with both primitives and complexes, lexical and compositional rules, different semantic types, functions combining with arguments, compositional determination of whole-meanings from part meanings, and so on.[21] What is rejected is just the idea that the primitives stand for real world objects and sets outside the mind, and that sentences express world-bound thoughts (relative to a set of parameters). Thus, the two objections from ontology in no way force one to abandon the discipline of semantics, or its many results. (We will revisit the tasks of semantics at the end of the chapter, when discussing the positive alternative to reference-based semantics.)

36.1.3 The Failure of Compositional Referential Semantics

As presented here, the former two arguments have focused primarily on public *words* (e.g. 'Harvard', 'photographer') and external *objects* (e.g. London). But the Chomskian target is the whole tradition of treating natural languages as, in key respects, like the formal languages invented by logicians — and, as explained above, that tradition brings with it a view not just about words and objects, but also views about predicates (e.g. verbs and adjectives) and sets, and sentences and truth. Traditionally, sentences are assigned "proto-thoughts": things which, given a particular setting for a fixed cluster of parameters (time, place, speaker, addressee), are true or false. And predicates are assigned sets of objects as their extension. These ideas are equally question-able, says the Chomskian — even putting aside the issues, raised above, of whether there *are* "public sentences" and "public predicates", and whether sets of common sense objects and proto-thoughts can be "seen" from the scientific perspective.

As noted at the outset, the mark of a Chomskian, as I here intend the term, is a commitment to methodological naturalism. Now, the methodological naturalist,

---

[21] See Jackendoff, 1983, **1991, 1996,** 2002 **for detailed examples of keeping much of this structure, but without external reference based semantics. It is worth stressing, however, that some of the authors discussed here remain highly skeptical about retaining this traditional superstructure, within a naturalist framework. See in particular Moravcsik 1998, who maintains that Fregean and Tarskian systems (i) need sharply defined word meanings, not permitting polysemy, and (ii) require that syntax mirror semantics — neither of which conditions, Moravcsik argues, hold for natural languages.**

inquiring into language as an aspect of nature, will follow the canons of the sciences, and **will** seek out, as her object of inquiry, a real object—possibly differing in important ways from pre-theoretical conceptions — that is scientifically tractable. That object turned out to be the language faculty: that innately provided, specialized module of the mind–brain. For this reason, the methodological naturalist will approach semantics, like phonology and syntax, as part of the study of the language faculty. What will emerge below, however, is that it is not plausible that the language faculty, taken alone, can assign proto-thoughts to sentences, or sets of objects to verbs, adjectives, etc. Hence sentence meanings cannot be thoughts, nor even "thoughts relative to a set of parameters"; and predicate meanings cannot be sets of objects in the external world.[22]

Let's begin with sentences. The fundamental points here are made by Pietroski (2003, 2005), building on Chomsky (1977). On the one hand, there is no *empirical* reason for thinking that what the language faculty assigns to a sentence will be capable of being true or false, even given contextual parameters like time, place, speaker, hearer, etc. (There's lots of empirical reason for thinking that *people* can say, and think, things that are true or false; but that is another matter.) The only thing which drives one to this expectation is, at bottom, a dubious analogy between natural objects and artifacts whose properties are stipulated (e.g. the predicate calculus). For the methodological naturalist, that in itself is damning. On the other hand, there is lots of empirical evidence that the language faculty alone doesn't assign thoughts (or propositions, or truth conditions, or what have you). In particular, very many sentences either lack truth conditions altogether, or are assigned truth conditions only via the rich interaction of different mental faculties.

Consider, for instance, the following pair of sentences:

(1) Poems are written by fools like me
(2) Mountains are climbed by fools like me

Putting aside the difficult question of what a referential semantics would assign as meanings for the parts (e.g. what real-world object does the plural word 'poems' stand for?), it does seem that the same *kind* of meaning, whatever it is, would have to be assigned by the language faculty to 'poems' and 'mountains', and to 'are written' and 'are climbed'. Moreover, the same syntactic structure appears in both sentences. Thus, the prediction would be that, as far as the language faculty goes, (1) says about poems and being-written whatever (2) says about mountains and being-climbed. Yet, insofar as they assign a proto-thought at all, the proto-thought that an agent would typically associate with (1) requires that *all* poems are written by fools, while

---

[22] Looked at slightly differently, as Fodor (2001) has argued, if semantics has to yield "a thought expressed (give or take a bit)", then a compositional semantics for natural language is just not possible. What one should conclude, says the Chomskian, is not that natural languages lack compositional semantics, which is Fodor's conclusion, but rather that scientific semantics just shouldn't be in the business of assigning thoughts to sentences—nor even "thoughts give or take a bit". Instead, the linguistic meanings of whole sentences just are those thiigs—whatever they turn out to be—which are compositionally determined from part-meanings plus syntax.

the proto-thought that an agent would typically associate with (2) does not require, for its truth, that *all* mountains are climbed by fools. Thus the truth conditions that language users tend to assign are importantly different. Nor is this an isolated example. Think, for instance, of how knowledge of the world impacts on what sense one assigns to 'may' in 'Marta may get cancer' versus 'Marta may smile if she wishes to.' In so far as one treats these as truth evaluable at all, one hears the first 'may' as expressing (epistemic or physical) *possibility,* and the latter 'may' as expressing *permissibility.* Moreover, this contrast in how 'may' is understood arises because, as a matter of fact, one isn't given permission to get cancer; and because it's too obvious to bear mention that it's (physically and epistemically) possible for a person to smile, if she wishes to. Or again, what the concatenation of a nominal and modifier contributes to meaning varies widely from case to case; in particular, the meaning of the nominal-modifier complex frequently reflects facts known about the world. For instance, compare 'Christmas cookie' ("made to be consumed at"), 'Girl guide cookie' ("sold by"), 'oatmeal cookie' ("made of"), 'yellow cookie' ("coloured"), 'fortune cookie' ("containing"), 'doggie cookie' ("made to be eaten by"), and 'Walmart cookie' ("sold at"). Given that concatenation doesn't always mean the same thing, if an agent is able to assign truth conditions at all to 'Phyllis ate a Boy Scout cookie', it is because the agent knows the relevant facts *about Boy Scouts:* e.g. that they are not an appropriate ingredient for a cookie, that there isn't a Boy Scout holiday, that Boy Scouts are too big to be inside a cookie, etc.

One might reply that language users are simply poor judges of what sentences mean: "the folk" mix up what the expression means with how it would be stand-ardly used. For example, while language users do indeed assign quantificationally different truth conditions to the sentence types (1) and (2), they are wrong to do so. This is a fair point. Indeed, Chomskians themselves are wont to note that lan-guage users do not have direct insight into the syntax of their own idiolect: language users are apt on first hearing, for instance, to incorrectly judge that 'The horse raced past the barn fell' is ungrammatical. Still, if we ask the semantic theorist what the proto-thought expressed by (1) *is,* since by hypothesis it isn't that *all* poems are writ-ten by fools like me, it will become clear that the meaning that the language faculty taken alone assigns is not the sort of thing which, even relative to a set of parameters (speaker, addressee, time, place), is true or false. Similarly for 'Marta may get can-cer' and 'Phyllis ate a Boy Scout cookie.' Whatever the context-invariant meaning of these sentences is, it's something *much* more abstract than a thought, or even a function from a restricted set of parameters to a thought.[23] Indeed, this attempted reply brings out even more clearly the deeper problem with trying to assign truth conditions to sentences. That deeper problem has to do with where the differences in truth conditions that agents assign—e.g. between (1) and (2)—derive from. Part

---

[23] The point is, of course, closely related to the idea—defended by Robyn Carston, Francois Récanati, John Searle, Dan Sperber and Deirdre Wilson, and Charles Travis—that there are "pragmatic determinants" of the truth conditions of speech acts. See Carston 2002 for detailed and illuminating discussion.

of the Chomskian point is precisely that the difference in the truth conditions that language users assign, in so far as they do assign truth conditions, derives from real world knowledge that people have—e.g. about poems versus mountains (i.e. that the former are all human creations, but the latter are not). Moreover, there is no other source of *truth* conditions: leave that real-world knowledge out and what is determined is too abstract to bear a truth value. Thus, in so far as we agents assign truth conditions to sentences at all, the truth conditions we assign are a massive interaction effect of different kinds of knowledge: knowledge afforded by the language faculty, yes, but also knowledge afforded by many other parts of the mind–brain. Not being solely an aspect of the language faculty, it follows that the truth conditions which people assign to sentences do not fall within the domain of the science of language. (See Borg, this volume, for a rather different view.)

Another kind of case raises problems for the idea that predicates (verbs, verb phrases, adjectives, adjective phrases, etc.) have sets of external objects as their content. Compare sentences (3) through *(6):*

(3) The house is green
(4) The ink is green
(5) The banana is green
(6) The stoplight is green

In each case, in so far as talk of "contributing sets" is appropriate at all, [₁′ is green] appears to be contributing a quite different set in the four cases. In (3), the house must be in the set of things which are green on the outside (though the house need not be entirely green on the outside). Similarly for (5), which requires only that the banana peel be green. In contrast (4) requires that the stuff (which right now looks black) be in the set of things which, when applied to paper and allowed to dry, will be green. As for (6), the science of colour tells us that the property exhibited by the stoplight is physically very different from that exhibited by the banana peel in (5): the stoplight being green involves not the reflection of light, but the emission of light. So, thought of as a physical set, the one which [₁is green] picks out in *(6)* is very different yet again. (See: Moravcsik 1998: 44–5 for similar remarks about 'is white'; Jackendoff, 1991: 44 on different senses for 'end'; and Jackendoff, 1983, 2002 and Pustejovsky, 1995 for a panoply of other examples.) *Part* of what appears to be going on here is this: which set of things is associated by speakers with one part of the sentence depends upon what they associate with the other parts. Here, the set that speakers associate with [₁′ is green] depends on the kind of thing that 'green' is thought of as applied to: houses, ink, bananas, etc. Worse, the variation in the set selected by the agent as the denotation for 'green', depends upon facts about how reasonable speakers would use the sentence—which in turn depends on factors like how likely it is for houses to be wholly and completely green, what ink is used for, etc. Once again, then, the denotation, in so far as language users assign one at all, is an enormous interaction effect, and does not depend solely upon the language faculty; so, there can be no science of language which assigns sets to predicates.

One natural reaction to this kind of example, an idea pursued by Jerry Fodor and Ernie Lepore in a series of articles, is to say that 'green' simply means *green* in (3) through (6) — adding that there are lots of different ways for things to be green. (See especially Fodor and Lepore, 1998.) The point is well taken. But green, so construed, now ceases to be a mind-independent property "out there in the world", to which words may simply attach. Of particular importance here, the set univocally denoted by 'green' becomes a set that no proper science could treat of — precisely because that collection of objects becomes wildly heterogeneous from a scientific perspective. What the house, the ink, the banana and the stop light have in common, in being green, is not something that any genuine science can see. To put the point differently, note that one might equally claim — and it isn't exactly false — that 'in' just means in as it occurs in 'a boy in trouble', 'a hole in her sock', 'a flaw in my argument', 'a detective in the novel', 'a C-sharp in the symphony', and so on. Even allowing that 'in' always means in, it clearly won't follow that there is scientifically tractable thing, "in-ness", denoted in all these cases. The worry is that the same holds for 'green', and the "green-set" univocally picked out by this word: if *that* is what 'green' stands for, we get a "single constant referent", but we cannot have an explanatory science that describes the word–world relation in the case of 'green'.

Once again, then, we have a reason for expecting that the science of language — which, for the methodological naturalist, is about the language faculty — cannot, and should not try to, assign as meanings the kind of thing that gets assigned in invented logical languages.

Actually, some would draw a stronger conclusion than the one I have been defending at length, viz. that a comprehensive *science of language* won't treat of word–world relations. Some might additionally conclude that, being a massive interaction effect of different causes, *no* genuine science will take the truth conditions we assign to sentences (or the sets we assign to predicates), as a thing to be explained — since genuine sciences are in the business of describing causal forces, not such highly complex particular effects.[24] To offer a comparison that Chomsky himself often gives, it is not the business of any science to describe the trajectory of a given falling leaf — even though it's quite true that scientific laws together contribute to how the leaf in fact fell. (Actually, there are two reasons why no science describes the falling of an individual leaf: first, it is non-tractable; second, it's just not interesting. I presume that the problem about any science assigning truth-conditions to sentences is the lack of tractability, not a lack of interest. But it might be both.) This may seem to go too far: surely it's altogether implausible that no science can capture such interaction effects. Part of the implausibility fades, however, when it's recalled that not every empirical enterprise that attempts to systematize is a genuine science; or anyway, it's not a science in the sense intended here. In the sense intended here, science involves explaining seemingly simple phenomena by postulating unobservables; and it involves the aim of

[24] Some of the authors listed at the outset explicitly disavow the idea that science, even "genuine explanatory science", is limited in this way. In particular, both Jackendoff (p.c.) and Moravcsik (1998, 2002) are much more sanguine than Chomsky is, about genuine sciences of complex interactions, including sciences of human creations.

integration with other "core natural sciences". It is, then, an open question whether the so-called "sciences of the complex", not to mention the "social sciences", really are *sciences* in the sense in question. No doubt the former uncover statistical patterns in the weather, ecosystems and the stock market, using sophisticated empirical methods; and the latter undoubtedly state empirically discovered rough generalizations about cultural products. And maybe similar techniques could be applied to the assignment of truth and satisfaction conditions. But this isn't enough. (To come at the point another way, it wouldn't be unreasonable to say that genuine natural science, in the sense in question here, is such a special enterprise, that it is a relatively recent arrival on the scene: it simply did not exist before the Renaissance. The question, then, is whether *that* "special enterprise" can be applied to massive interaction effects. Put this way, the extra worry, that goes beyond the main conclusion argued for in the article, is that this special enterprise cannot be so applied, so that, as a case in point, no science (of language, or of anything else) can treat of truth conditions for sentences or satisfaction conditions for predicates.)

36.2 Remarks on a Positive Alternative

So, what, from a methodological naturalist perspective, *does* the meaning of expression look like, if it isn't a matter of a word–world relationship? Well, meaning looks a lot like syntax. In particular, it looks like syntax *which has procedural implications.* Indeed, it isn't too far wrong to say that meaning is that aspect of natural language syntax which plays a causal role in the conceptual-intentional system (and ultimately, though in very unclear ways, in thought, and in the production of behavior). Though rather short on details, this broadly Chomskian idea can be fleshed out a little by making a comparison with phonetic features. Phonetic features are, for Chomsky and many other linguists, mental instructions, which are hooked up, in the first instance, with other representational systems — where something counts as an "instruction" because of its intricate form-based causal powers, not because it is contentful in the sense of standing for something in the external world; and where 'representation' is stripped of its philosophical "standing for" connotations. These latter representational systems ultimately contribute to moving tongue, lips, etc., thereby playing a crucial part in giving rise to speech sounds.[25] In the same sort of way, the meaning of an expression, on this approach, is a cluster of semantic features that similarly interface with (a rather different) mental system, the "conceptual and intentional system". And this and other systems play a part in actions by the agent. Chomsky writes:

---

[25] As Jackendoff, 2002, points out, these features are also instructions from the auditory system; and that system presumably cannot output motor instructions. So, identifying phonetic features with instructions to the sensorimotor system is clearly an oversimplification.

Each expression can be regarded as a collection of information for other systems of the mind–brain. The traditional assumption, back to Aristotle, is that the information falls into two categories, phonetic and semantic; information used, respectively, by sensorimotor systems and conceptual-intentionalsystems —the latter "systems of thought," to give a name to something poorly understood. (2002: 87)

(See Chomsky, 1997 for further discussion, and also Chomsky, 2000c: 90–1, where the notion of "instruction" in particular is discussed in a bit more detail.)

Crucially, semantics in this tradition can be nothing more than rules for mapping one mental representation to another, by well-defined tractable procedures. The science of language is thus restricted to describing the sub-personal, unconscious, automatic, cognitively impenetrable rules of the language faculty. Put in a nutshell, it is restricted to this because only this is formally tractable. The personal-level, conscious, reasoned and flexible *use* of language, to talk about the world, is excluded from the domain of science, properly so called.

Put this way, it can seem that semantics becomes extremely "thin", so that not much can be said about it. But that would be a mistaken impression. First, as McGilvray (1998) stresses, internalist semantics still faces the enormous task of finding out what the various "meaning features" are, and finding out which lexical items exhibit which features. (An especially nice example of how this task is pursued is Jackendoff, 1991.) Given the centrality of feature checking to Minimalism, hinted at above, lexical semantics of this kind is a very important task indeed. But there is also lots of work to do on the "compositional" side of semantics.

To give the flavor of how the semantics of syntactic complexes proceeds, consider two examples. (For many other early examples, see Jackendoff, 1983.) Compositional semantics, as reconceived, will still need to explain why, for example, whereas (7) is ambiguous (it can mean both *You want who to shoot?* and *You want to shoot who?*), sentence (8), with 'want' and 'to' contracted into 'wanna', is not ambiguous, and can only mean *You want to shoot who?*

(7) Who do you want to shoot?

(8) Who do you wanna shoot?

A partial explanation of this meaning-fact, simplifying greatly, goes as follows. Underlying (7) there are actually *two* syntactic structures, namely (9) and (10):

(9) Who_1 do you want t_1 to shoot?

(10) Who_1 do you want to shoot t_1

In contrast, only (10) underlies (8), because the trace t_1 in between 'want' and 'to shoot' in (9) blocks contraction: 'want t_1 to shoot' cannot become 'wanna shoot'. That is why (7) is ambiguous, but (8) is not. What remains to be said, to explain why (7) and (8) have the meaning they do, is to sort out why the two structures (9) and (10) mean what *they* do. This depends upon what the words mean, of course, which is the same in the two cases. It also depends upon what a trace co-indexed with 'who' contributes to meaning. It is this latter contribution which is different in (9) and (10), because of where the trace appears: because of the trace, in one case the direct object position is queried, in the other case it is the subject position that is queried. What's

important for present purposes is not whether this explanation is correct, or complete; what matters is that specifying all of these things remains the job of semantics, even once external world reference and "proto-thoughts" are put aside.

To give another familiar kind of example, in (11) 'him' can be referentially dependent on 'Juan' (that is, put crudely, 'him' is allowed to, though it need not, take its meaning over from the name); but in (12) 'him' *cannot* be referentially dependent on 'Juan'.

(11) $Juan_1$ asked Maria to kill him_1 [Juan = him, is possible]
(12) *$Juan_1$ promised Maria to kill him_1 [Juan ≠ him]

This is a phenomenon that needs to be explained. A by-now traditional explanation, again simplifying for present purposes, goes like this.[26] First, the underlying structure of the two expressions is a bit more complex than what appears in (11) and (12). At a minimum, we need to add an unpronounced subject PRO for the embedded infinitival clause 'to kill him', and we need explicitly to bracket off this embedded clause:

(13) $Juan_1$ asked $Maria_2$! [$_S$ PRO_2 to kill him_1]   [Juan = him, is possible]
(14) $Juan_1$ promised $Maria_2$! [$_S$ PRO_1 to kill him_1]   [Juan ≠ him]

Now, continues the story, it is a semantic feature of the verb 'promise' that its subject gets co-indexed with the *subject* of the embedded clause that follows, here the sentence 'PRO to kill him'. Because of this lexical semantic fact about 'promise', the PRO subject of the embedded clause [$_S$ PRO to kill him_1], in (14), comes to share the index 1 both with 'Juan' and with 'him'. (This contrasts with 'ask', which is a verb whose *object* gets co-indexed with the subject of the embedded clause, as in (13); that is why the PRO subject of (13) shares the index 2 with 'Maria'.) But, as a result of a general restriction that needn't detain us here, in a simple sentence of the form 'SUBJECT kill him', 'him' *cannot* be referentially dependent on the subject phrase.[27] To see the pattern, note that 'him' cannot be referentially dependent on 'Juan', 'The man' or 'He' in (15)–(17):

(15) *$Juan_1$ killed him_1 [Juan ≠ him]
(16) *The man_1 killed him_1 [The man ≠ him]
(17) *He_1 killed him_1 [He ≠ him]

Given the semantic properties of 'promise', the co-reference principle that underlies (15)-(17), and the postulated element PRO in the embedded infinitival clause, the semantic fact that (12) cannot mean *Juan promised Maria to kill himself* is now partially explained. Again, what really matters for present purposes is not whether this

---

[26] In both of the compositional examples discussed here, the principles I appeal to are now thought to follow from deeper constraints. Indeed, in Minimalism all "rules" end up being typological artifacts of (i) lexical features, (ii) some very basic operations (e.g. Merge, Agree), (iii) overarching economy conditions (e.g. simplicity and locality), and (iv) output conditions imposed by the two interfaces. See Chomsky, 1995a, 2000c, 2001. I employ the older framework, however, because explanations in those terms are rather easier to present, and they exemplify equally well the kind of task that remains, even after reference-based semantics is abandoned.

[27] The general principle is Principle B of the Binding Theory: "A pronominal is free in its governing category" (Chomsky, 1981: 188).

explanation is precisely the right one; for present purposes, the key point is simply that this explanation draws on facts about the contrasting lexical semantic features of 'ask' and 'promise', on facts about structural constraints on co-indexing, and on facts about what co-indexing contributes to meaning. Here again, these are semantic issues that do not simply melt away with the rejection of reference-based semantics.[28]

In sum, as Pietroski (2005) concludes, "Trading in truth-values (and entities referred to)... does not change the basic questions. We still want to know, for any given sentence: what is its structure; what does it mean; and how is the former related to the latter?" Thus, semantics remains rich (and central to linguistic theorizing), and the sub-discipline of semantics still has much work to do in capturing the semantics properties of expressions — primitive and complex.[29]

But, it will be asked, if the science of language cannot ascribe real-world referents to words, and instead merely pairs linguistic representations with other linguistic expressions, and with inner representations of other kinds, how on earth does our talk manage to be about the world outside us? This is an exceedingly difficult question. The short answer is that, though words themselves don't refer,[30] *people* can refer using them. Nothing said above rules this out. Our speech acts and our thoughts are about the world — but not because of a relationship between particular natural language representations and particular outer things.[31] The long answer is.. . Well, no one knows what the long answer is.

In light of the short answer, one might hold out hope that there could be a science of *speaker reference.* Chomskians aren't optimistic about that, however: there cannot, they think, be a science that captures episodes of people referring either. First, speaker-reference is as much a massive interaction effect as speech episodes are in general — which entity the speaker manages to refer to, using 'he' or 'The woman from Spain', will clearly depend on a host of things. (Just as whether an expression "sounds right" will depend on many, many things beyond what the grammar states about the expression.) But, as hinted above, because of intractability (and sometimes because it's uninteresting), it may be that genuine sciences aren't in the business of

---

[28] Moreover, semantics as reconceived here will still explore relations *between* expressions, noting (and trying to explain) what logical entailments hold on the basis of meaning alone, which expressions are and are not synonymous, etc. Thus, to give but one example, semantics will try to explain why 'Saima persuaded Moonisah to leave' entails 'Moonisah formed the intention to leave'.

[29] David Lewis (1970: 190) famously complained that "Semantics with no treatment of truth conditions is not semantics". This slogan seems to have exerted enormous influence in philosophy of language. But, as Pietroski (2005) argues at length, at bottom Lewis is simply stipulating a usage for a technical term, 'semantics'. The methodological naturalist will eschew such stipulations, and will instead look for a real feature of the world to study. See also Jackendoff (1987: ch. 7) for related points about "Lewis's terminological imperialism" (1987: 130), and an early and extensive defense of internalist semantics in the face of Lewis' criticisms.

[30] To be clear, Chomsky does allow for a notion of reference for expressions. He labels it "relation R . But relation R does not introduce a relation between external entities and words; it is thoroughly internalist. See Chomsky, 1992a: 39 for discussion.

[31] Chomsky (1993: 22) does note that the use of 'refer' as applied to words is a technical coinage. However, as is dear from what has been said above, his reasons for saying that people refer, rather than having words refer, are not based on this minor detail about ordinary usage.

describing effects deriving from such multiple and varied causes; rather, genuine sciences are in the business of abstracting away to the causal forces that produce these effects. Thus, just as there is no science of which things "sound right", and no science of Chomsky's falling leaf, there may well be no science of what the person, in this particular circumstance, refers to. At least not in the sense of 'science' in play here. Second, and deeper, in so far as referring is something that the whole agent does, it is a conscious act of free will and reason. And, for Chomsky and some of his followers, that in itself puts it outside the scope of the sciences: for this reason alone, reference by speakers cannot be treated naturalistically either.;'

However you slice it, then, meanings just *are* in the head. Or anyway, there can be no comprehensive science of language which studies "meanings" of the word–world variety: that kind of meaning-theory just isn't scientifically tractable. As Chomsky (1992a: 45) succinctly puts the conclusion, "Naturalistic inquiry will always fall short of intentionality."

REFERENCES

Antony, Louise (1997)."Meaning and Semantic Knowledge". *Aristotelian Society Supplementary Volume*, 71: 177–209.

Bencivenga, Ernianno (1983). "An Epistemic Theory of Reference", *The Journal of Philosophy*, 80: 785–803.

Bilgrami, Akeel (2002). "Chomsky and Philosophy". *Mind and Language*, 17: 290–302.

Carston, Robyn (2002).*Thoughts and Utterances.* Oxford: Blackwell.

Chonisky, Noam (2002). *On Nature and Language.* Cambridge, UK: Cambridge University Press.

_____ (2001). "Derivation by Phase". In M. Kenstowicz (ed.) *Ken Hale: A Life in Language.* Cambridge, MA: MIT Press.

_____ (2000a). *New Horizons in the Study of Language and Mind.* Cambridge, UK: Cambridge University Press.

_____ (2000b). "Internalist Explorations". In Chonisky, 2000a.

_____ (2000c). "Minimalist Inquiries: The Framework". In R. Martin, D. Michaels, and J. Uriagereka (eds.) *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik.* Canibridge, MA: MIT Press.

_____ (1997)."New Horizons in the Study of Language". In Chonisky, 2000a.

_____ (1995a). *The Minimalist Program.* Cambridge, MA: The MIT Press.

_____ (1995b). "Language and Nature". *Mind*, 104: 1–61. Reprinted in Chomsky, 2000a.

_____ (1994). "Naturalism and Dualism in the Study of Mind and Language". *International Journal of Philosophical Studies*, 2: 181–200. Reprinted in Chomsky, 2000a.

_____ (1993). *Language and Thought.* London: Moyer Bell.

_____ (1992a). "Explaining Language Use". *Philosophical Topics,* 20: 205–31. Reprinted in Chomsky 2000a.

¹² Chomsky writes: "It is not excluded that human science-forming capacities simply do not extend to this domain [i.e. how stimulus conditions given a cognitive state give rise to behavior (including the use of language)], or any domain involving the exercise of will, so that for humans, these questions will always be shrouded in mystery" (1975: 25). Also: "The phrase 'at will' points to an area beyond serious empirical inquiry" (Chomsky, 2002: 59).

_____ (1992b)."Language and Interpretation: Philosophical Reflections and Empirical Inquiry". In J. Earman (ed.) *Inference, Explanation and Other Frustrations: Essays in the Philosophy of Science.* Berkeley: University of California Press, pp. 99–128. Reprinted in Chomsky, 2000a.

_____ (1986). *Knowledge of Language.* New York: Praeger.

_____ (1981). *Lectures on Government and Binding.* Dordrecht: Foris.

_____ (1977). *Essays on Form and Interpretation.* New York: North Holland.

_____ (1975). *Reflections on Language.* New York: Pantheon Books.

Davidson, Donald (1984). *Inquiries into Truth and Interpretation.* Oxford: Oxford University Press.

_____ (1976). "Reply to Foster". In G. Evans and J. McDowell (eds.) *Truth and Meaning.* Oxford: Clarendon Press, pp. 33–41. Reprinted in Davidson, 1984.

_____ (1967). "Truth and Meaning" *Synthese,* 17: 304–323. Reprinted in Davidson, 1984.

Fodor, Jerry (2001). "Language, Thought and Compositionality". *Mind & Language,* 16: 1–15.

_____ (1983). *Modularity of Mind.* Carnbridge, *MA:* MIT Press.

Fodor, Jerry & Ernie Lepore (2002). *The Compositionality Papers.* Oxford: Oxford University Press.

_____ (1998). "The Emptiness of the Lexicon: Reflections on Pustejovsky". *Linguistic Inquiry,* 29: 269–88. Reprinted in Fodor and Lepore, 2002.

Foster, John A. (1976). "Meaning and Truth Theory". In G. Evans and J. McDowell (eds.) *Truth and Meaning.* Oxford: Clarendon Press, pp. 1–32.

Hornstein, Norbert (1989). "Meaning and the Mental: The Problem of Semantics after Chomsky". In A. George (ed.) *Reflections on Chomsky.* Oxford: Blackwell, pp. 23–40.

_____ (1984). *Logic as Grammar.* Cambridge, *MA:* The MIT Press.

Jackendoff, Ray (2002). *Foundations of Language.* Oxford: Oxford University Press.

_____ (1996). "Semantics and Cognition". In S. Lappin (ed.) *The Handbook of Contemporary Semantic Theory.* Oxford: Blackwell, pp. 539–559.

_____ (1991). "Parts and Boundaries". *Cognition,* 41: 9–45.

_____ (1987). *Consciousness and the Computational Mind.* Cambridge, *MA:* The MIT Press.

_____ (1983). *Semantics and Cognition.* Cambridge, *MA:* The MIT Press.

Lewis, David (1970). "General Semantics". *Synthese,* 22: 18–67. Reprinted in *Philosophical Papers,* Volume 1. Oxford: Oxford University Press, pp. 189–229.

Pustejovsky, James (1995). *The Generative Lexicon.* Cambridge, *MA:* The MIT Press.

Matheson, David and Robert J. Stainton (2002). "Varieties of Empiricism". In Y. Bouchard (ed.) *Perspectives on Coherentism.* Aylmer, Quebec: Éditions du Scribe, pp. 99–113.

McGilvray, James (1999). *Chomsky: Language, Mind, and Politics.* Cambridge, UK: Polity Press.

_____ (1998). "Meanings Are Syntactically Individuated and Found in the Head" *Mind and Language,* 13: 225–80.

Montague, Richard (1974). *Formal Philosophy.* New Haven, CT: Yale University Press.

Moore, G. E. (1939). "Proof of an External World". *Proceedings of the British Academy,* 25: 273–300. Reprinted in T. Baldwin (ed.) (1993) G.E. *Moore: Selected Writings.* London: Routledge.

Moravcsik, Julius (2002). "Chomsky's New Horizons". *Mind and Language,* 17: 303–11.

_____ (1998). *Meaning, Creativity and The Partial Inscrutability of the Human Mind.* Stanford: CSLI.

Pietroski, Paul (2005). "Meaning before Truth". In G. Preyer and G. Peters (eds.) *Contextualism in Philosophy.* Oxford: Oxford University Press, pp. 255–302.

Pietroski, Paul (2003). "The Character of Natural Language Semantics". In A. Barber (ed.) *The Epistemology of Language*. Oxford: Oxford University Press, pp. 217–56.

Radford, Andrew (1997). *Syntactic Theory and the Structure of English*. Cambridge, UK: Cambridge University Press.

Smith, Barry C. (1992). "Understanding Language". *Proceedings of the Aristotelian Society*, 92: 109–41.

Smith, Neil (1999). *Chomsky: Ideas and Ideals*. Cambridge, UK: Cambridge University Press.

Strawson, Peter F. (1950). "On Referring" *Mind*, 59: 320–44. Reprinted in R. J. Stainton (ed.) *Perspectives in the Philosophy of Language*. Peterborough, ON: Broadview.

Wiggins, David (1997). "Languages as Social Objects". *Philosophy*, 72(282): 499–524.