

IO

Meaning before Truth

PAUL M. PIETROSKI

According to Chomsky (1996: 52),

we cannot assume that statements (let alone sentences) have truth conditions. At most, they have something more complex: ‘truth indications’, in some sense. The issue is not ‘open texture’ or ‘family resemblance’ in the Wittgensteinian sense. Nor does the conclusion lend any weight to the belief that semantics is ‘holistic’ in the Quinean sense that semantic properties are assigned to the whole array of words, not to each individually. Each of these familiar pictures of the nature of meaning seems partially correct, but only partially. There is good evidence that words have intrinsic properties of sound, form, and meaning; but also open texture, which allows their meanings to be extended and sharpened in certain ways; and also holistic properties that allow some mutual adjustment. The intrinsic properties suffice to establish certain formal relations among expressions, interpreted as rhyme, entailment, and in other ways by the performance systems . . .

If this is right, and I think it is, we must re-evaluate many widely accepted assumptions about meaning, truth, and context-sensitivity.

I Overview

Chomsky offers a plausible though often ignored conception of linguistic meaning and its relation to truth: the meaning of a natural language sentence *S* is an internalistic property of *S*, determined by the human language faculty and the relevant lexical items; the semantic properties of sentences, which reflect how human beings understand natural language, are theoretically tractable; but if an utterance of *S* is true or false, its truth or falsity is typically

254 COMPOSITIONALITY, MEANING, AND CONTEXT

a massive interaction effect due to the meaning of S and many factors not indicated by elements of S.¹ In my view, this conception is preferable to more standard alternatives, which either (i) burden theories of meaning with implausible predictions, or (ii) abandon good explanations.

Davidson (1967*a*, 1984) conjectured that there are Tarski-style theories of truth for natural languages, and that such theories can serve as theories of meaning. This proposal was very useful, but too bold. Sentences like (1–3) illustrate difficulties, discussed in later sections.

- (1) France is hexagonal, and it is a republic
- (2) This government does little for the sake of the average American, whose children will inherit the massive deficit that is accumulating
- (3) Hamlet lived with his parents in Denmark

For example, I don't think 'France' has the semantic correlate it would need to have, given a compositional theory of truth. But we shouldn't conclude that there are no theories of meaning for natural languages. We should conclude that such theories are not theories of truth. Correlatively, the meanings of declarative sentences do not specify truth-conditions, not even relative to contexts. In epistemic mode: knowing what sentence S means—that is, understanding S—is not a matter of somehow associating S with a function from contexts to truth-conditions.

Rather, the meaning of S is a compositionally determined intrinsic property of S that *constrains and guides without determining* how S can be *used to make* true or false assertions in various conversational situations. A related theme, often stressed by Chomsky, is that we should combine the idea that sentences have intrinsic semantic properties with a cluster of claims associated with Wittgenstein, Austin, and Strawson: making truth-evaluable assertions is *one* of the things we can do with sentences, in contexts, though uses of this kind are highly variable; while people refer to things, words don't; and sentence *use* may not be a theoretically tractable phenomenon. So if we adopt the good idea that theories of meaning are theories of *understanding*, we should not expect a tight connection between meaning and truth.

¹ See Pietroski (2003*a*) for discussion focusing mainly on Chomsky (1977, 2000). Let me note that what I say there, and here, is partly a result of many conversations with Jim McGilvray. For similar views with different emphases, see Moravcsik (1977, 1998), Hornstein (1984), McGilvray (1996, 1999), Hinzen (2002); see Stainton (forthcoming) for useful review and discussion. Many linguists may adopt some such view in practice, and regard their claims to the contrary as dispensable idealizations; but cf. Higginbotham (1989*a*, 1989*b*).

This conflicts with some Quinean/Davidsonian claims about the nature and source of semantic phenomena. But we should reject these claims in any case. We may eventually earn the right to say that semantic theories associate sentences with ‘truth indications’ in some interesting sense. Until then, it may be best to just say that theories of meaning/understanding for natural languages are *like* theories of truth for formal languages in certain specifiable respects. For this may be all we need, in order to explain what semantic theories actually explain—for example, facts about entailment relations, ambiguity, and how natural language *cannot* be understood.

In the next section, I lay out some spare assumptions about the enterprise of semantics. This precludes certain conceptions of meaning, given some observations due to Chomsky and others following him. These observations bolster the arguments, discussed below, for a Chomsky-style conception. Along the way, I briefly consider some alternatives. While my aim is not to establish that these alternatives are wrong, I do think they should be evaluated on their own merits, without assuming tendentious views about how meaning is related to truth.

2 Assumptions

Let’s say that a *natural* language is one that human children can acquire, in the normal course of development, given a course of experience that is not atypical for members of the relevant linguistic community. For these purposes, I take as given that a theory of meaning for a natural language is a theory of understanding, and thus a theory of certain human capacities; see Chomsky (1965, 1986), Dummett (1975), Higginbotham (1985). Such a theory can take the form of an algorithm that associates signals of the relevant language (sounds, in the case of a spoken language) with interpretations, leaving it open for now what interpretations are. But a theory of meaning for a natural language *L* is not merely an algorithm of this sort. It also purports to explain, at least in part, *how* a speaker of *L* associates signals of *L* with interpretations.²

² Dummett (1975, 1976) makes the important observation that a theory of understanding might have the *formal character* of a Tarski-style truth-theory without *being* a theory of truth. One can offer such a theory without construing the labels for “truth values” in terms of classical truth and falsity. But my point will not be that we should replace ‘is true’ with ‘is assertable’.

256 COMPOSITIONALITY, MEANING, AND CONTEXT

So if certain phenomena—ambiguity, entailment, indexicality, or whatever—reflect the way that speakers associate linguistic signals with interpretations, then these phenomena bear on theories of meaning. I do not deny that the right level of abstraction, whatever that turns out to be, will let us ignore many details about how particular speakers make such associations. One can understand my idiolect of English without biochemically associating signals with interpretations in the *exactly* the way I do. But one cannot stipulate, in advance of inquiry, which facts are (not) relevant for theories about how speakers of a language understand that language. Our job as theorists is to describe and explain the relevant facts, whatever those turn out to be.

2.1 Negative Facts are Relevant

Chomsky (1957, 1965, 1970, 1977, 1981, 1986) drew our attention the fact that for any signal σ of a natural language, there are endlessly many interpretations that σ cannot have, and that *nonambiguity* often calls for explanation. Consider the contrast between (4) and (5),

- (4) John is easy to please
- (5) John is eager to please

which can be paraphrased with (4*a*) and (5*a*), but *not* with (4*b*) and (5*b*).

- (4*a*) It is easy for us to please John
- (4*b*) #It is easy for John to please us
- (5*a*) John is eager that he please us
- (5*b*) #John is eager for us to please him

Every adult speaker of English knows what (4) and (5) mean, *and* what they don't mean. So evidently, if any normal human child undergoes any ordinary course of experience in any English-speaking community, that child will acquire an idiolect according to which the sounds of (4) and (5) are associated with the interpretations indicated with (4*a*) and (5*a*) *but not* the interpretations indicated with (4*b*) and (5*b*). This is an interesting fact. For there is no general prohibition against ambiguity in natural language.

Lexical ambiguity is ubiquitous. The sound of 'bear' can be associated with more than one interpretation (and spelling). So if expressions are signal–interpretation pairs, there are homophonous but distinct lexical expressions. Different expressions composed from the same (overt) lexical items can also be homophonous, as illustrated with (6) and (7).

- (6) The goose is ready to eat
 (7) The millionaire called the senator from Texas

The words in (6) can be combined to form a sentence meaning that the goose is prepared to dine, or a sentence meaning that the goose is *pret-a-manger*; compare ‘The goose is eager/easy to eat’. Similarly, (7) can be associated with the interpretations indicated with (7*a*) and (7*b*).

- (7*a*) The millionaire called the senator, and the senator is from Texas
 (7*b*) The millionaire called the senator, and the call was from Texas

But again, there is a negative fact, since (7) cannot have the interpretation indicated with (7*c*).

- (7*c*) #The millionaire called the senator, and the millionaire is from Texas

So it seems that any “Englished” child—that is, any normal human child who undergoes an ordinary course of experience in an English-speaking community—will acquire an idiolect according to which: (7) has a reading on which it implies (7*a*), and a reading on which it implies (7*b*), but no reading on which it implies (7*c*). More generally, if σ is a signal of a natural language L in which σ is associated with certain interpretations $\langle \mu_1 \dots \mu_n \rangle$ but not others, then (other things equal) an L -ed child will come to associate signals with interpretations in a way that associates σ with $\langle \mu_1 \dots \mu_n \rangle$ but not other interpretations; and this is so even for the endlessly many signals that the child will never encounter, at least not prior to meeting a linguist.

This bears on theories of meaning, since facts about how humans *don't* associate signals with interpretations may well reveal important aspects of *how* humans understand language—especially if such facts raise theoretically interesting questions about how children manage to converge (in so far as they do converge) on agreement about signal–interpretation associations, despite disparate and often relatively impoverished experience. And once it is granted that the explananda for semantics need not be limited to facts about what signals *do* mean, it quickly becomes clear that these “positive” facts reflect the tip of an iceberg; see Higginbotham (1985). Consider, as another illustration, the much discussed facts concerning the (im)possibilities for antecedence of pronouns in (8)–(10); see Chomsky (1981, 1986).

- (8) Pat thinks that Chris likes him/himself
 (9) Pat wants to meet Chris and like him/himself
 (10) Pat wants Chris to meet and like him/himself

258 COMPOSITIONALITY, MEANING, AND CONTEXT

In such cases, the unavailable interpretations do not correspond to incoherent or contradictory thoughts. And there are endlessly many word-strings for which an unavailable interpretation would be a *more* reasonable guess about what the speaker might have meant, compared with the mandatory interpretation (which may be initially hard to discern). Consider ‘Pat wants Chris to meet and like himself’ or ‘Was the hiker who lost kept walking’. Such examples also constitute data for “poverty of stimulus” arguments, according to which humans impose arational constraints on the space of possible interpretations for linguistic signals.³

Nonambiguity is intimately connected with entailment. If one sentence follows from another, that is an interesting fact, because distinct sentences are not typically related in this way. For example, we learn something about how humans understand natural languages by trying—as Davidson (1967*b*, 1985) did—to explain the facts illustrated with (11)–(13).

- (11) Brutus stabbed Caesar
- (12) Brutus stabbed Caesar with a knife
- (13) Brutus stabbed Caesar, and he did it with a knife.

The relevant explananda include the fact that (11) follows from (12) *and not conversely*, just as (11) follows from (13) *and not conversely*. Correspondingly, (12) can be paraphrased with (13) but not with (14), which follows from (11).

- (14) Brutus stabbed Caesar, or something was done with a knife

Similar points apply to (4)–(10). In most dialects, the sound of ‘Pat wants to meet Chris and like himself’ does not have a natural interpretation on which it follows that Pat wants to like Chris. We want to know why not. For whatever keeps speakers from hearing ‘himself’ as linked to ‘Chris’ in this example is potentially relevant for theories of meaning/understanding.

2.2 Disquotation is Inadequate

Horwich (1997, 1998) outlines a conception of meaning heavily biased towards positive facts. And while few other theorists adopt Horwich’s “deflationary” view in its entirety, I suspect that many are inclined to adopt something like it with regard to certain aspects of linguistic meaning. So

³ For reviews of some relevant literature, including psycholinguistic studies of young children, see Crain and Pietroski (2001, 2002); see also Hornstein and Lightfoot (1981), Laurence and Margolis (2001).

it is worth being clear about the kinds of explanation that get lost on such a view.

Horwich begins with a plausible idea: the meaning of a complex expression *E* is a property of *E* that can be viewed as the result of combining the meanings of *E*'s constituents in ways corresponding to ways (exhibited by *E*) of combining expressions. And he rightly notes that given this conception of meaning, accounting for the mere compositionality of semantic properties is not hard. In particular, one need not adopt a substantive theory of truth to say that the meaning of *E* is determined by (i) the meanings of *E*'s basic constituents, and (ii) the relevant ways of combining those lexical meanings. But Horwich also says that "Understanding one of one's own complex expressions (nonidiomatically) is, by definition, nothing over and above understanding its parts and knowing how they are combined" (1997: 504). On this view, if "one has worked out how a certain sentence is constructed from primitive syntactic elements", then "provided one knows the meanings of those elements" one understands the sentence "automatically and without further ado . . . no further process needs to be involved, leading from these initial conditions to the state of understanding the sentence".

Horwich concludes that, given the grammatical structure of a complex expression, all a "theory" of meaning/understanding must provide is a specification of what the lexical items mean; where such a specification can be given disquotationally, using "axioms" like '*barked* means BARKED' to report that a certain word means what it does. But this seems wrong, for reasons discussed by Higginbotham (1985) and others. To understand a complex expression *E*, one must also know how the form of *E* contributes to the meaning of *E*; and this imposes substantive constraints, not captured by disquotation, on what lexical items can(not) mean.

Recall (7), which has no reading on which it implies (7*c*).

(7) The millionaire called the senator from Texas

(7*c*) #The millionaire called the senator, and the millionaire is from Texas

As a matter of natural language *grammatical* structure, 'from Texas' cannot modify 'millionaire' in (7). But let's assume, if only for illustration, that it can modify 'senator' or 'called the senator' as shown in the homophonous (7 α) and (7 β).

(7 α) [[The millionaire][called [the [senator [from Texas]]]]]

(7 β) [[The millionaire][[called [the senator]][from Texas]]]

260 COMPOSITIONALITY, MEANING, AND CONTEXT

This accounts for the ambiguity of (7). But it does not explain why (7) has no reading on which it implies (7c). We need some further theoretical claims according to which the structure indicated in (7β) does *not* support the following interpretation: the millionaire is *both* someone who called the senator *and* someone who is from Texas. Of course, the structure doesn't support this interpretation. Speakers know, and so children somehow figure out, that (7β) is not associated with the interpretation indicated with (7c). But this *isn't* just a matter of "working out" how (7) can be constructed from the relevant words, and knowing what the words mean. One must also associate each way of structuring the words in (7) with the right interpretation, while *not* associating (7β)—or whatever the relevant grammatical structure is—with the interpretation indicated in (7c). And this is not trivial.

Davidson's (1967*b*) event analysis suggests a hypothesis. On any reading, (7) is understood partly in terms of an event variable associated with the verb 'call'; the grammatical subject and object are understood as indicating a *caller* and a *callee*; and in (7β), 'from Texas' is understood as a predicate linked to this variable, *but not* to potential callers. One can spell this out in several ways. Perhaps 'called' and 'called the senator' are understood as predicates associated with an event variable, *and not* with a variable corresponding to potential callers, while combining 'called the senator' with 'from Texas' signifies predicate-conjunction. On this view, the meaning of (7β) might be represented as follows: $\exists e[\text{Agent}(e, \text{the millionaire}) \ \& \ \text{Past-Calling}(e) \ \& \ \text{Theme}(e, \text{the senator}) \ \& \ \text{From}(e, \text{Texas})]$; where 'Past-Calling(*e*)' means that *e* was an event of calling. Alternatively, perhaps 'called' and 'called the senator' *are* associated with a variable for callers, but for some reason, 'from Texas' cannot be linked to this variable in (7β). On this view, the meaning of (7β) might be represented as follows: $\exists e[\text{Called}(e, \text{the millionaire}, \text{the senator}) \ \& \ \text{From}(e, \text{Texas})]$; where 'Called(*e*, *x*, *y*)' means that *e* was a calling by *x* of *y*.⁴

For present purposes, the details do not matter. The important point is that the relevant negative fact, concerning the interpretation of 'from Texas' in (7β), is not even adequately described—much less explained—with a semantic "theory" that simply provides a disquotational algorithm for associating signals (relative to grammatical structures) with interpretations.

⁴ For defense of the view that verbs like 'called' are understood as monadic predicates, whose grammatical arguments are associated with thematic roles, see Parsons (1990), Schein (1993, 2002, forthcoming), Pietroski (1998, 2002, 2003*c*, 2004), Herburger (2000); see Kratzer (1996) for a slightly different view that would have the same consequences for (7). For defense of the second view, see Higginbotham (1985); see also Taylor (1985).

One can say that just as ‘called’ and ‘Texas’ mean what they mean, so the grammatical relation corresponding to adverbial modification makes whatever semantic contribution it makes. But this does not explain why [[called [the senator]][from Texas]] cannot be understood as a predicate satisfied by x iff x called the senator and x is from Texas, or why (7) has a reading on which it implies that the *call* was from Texas. Explaining this requires substantive hypotheses about ‘called’, and the *specific* significance of combining expressions in certain ways.⁵ Likewise, we want a theory that explains relevant negative facts concerning ‘John is easy/eager to please’. And this will presumably require substantive hypotheses about what ‘easy’ and ‘eager’ mean.

One way or another, we need to capture the following idea: the meaning of ‘easy’ is lexicalized so that when this word combines with ‘to please’ and ‘John’, constraints on grammatical structure and compositional semantics conspire to ensure that John is said to be an individual who is easily *pleased*; while the meaning of ‘eager’ is lexicalized so that when this word combines with ‘to please’ and ‘John’, John is said to be an individual who is eager to be a *pleaser*. We want to know more about these facts, which seem to be symptoms of how lexicalization interacts with (syntactic and semantic) composition in natural language. But just saying that ‘easy’ has the semantic properties that it has, or that ‘eager’ applies to what it applies to, tells us nothing about how ‘easy’ and ‘eager’ differ in a way that ‘easy’ and ‘hard’ do not. And recall that ‘The goose is ready to eat’ is ambiguous; see Chomsky (1965, 1977, 1986). Thus, disquotational description blurs interesting distinctions, and it obscures relevant explananda.

In what follows, I assume that this is the normal case: interesting phenomena—relevant to theories of meaning, since they bear on linguistic understanding—are often due to subtle interactions between lexical items and natural composition; and explaining such phenomena typically requires substantive (non-disquotational) hypotheses about lexical meanings and

⁵ We could, after all, invent a language in which (7 β) has the meaning indicated with (7 c). And a Horwich-style theory would apply just as well to such a language. It is also worth noting that small grammatical differences can have significant semantic effects. In ‘I heard Pat sang’, ‘Pat’ occupies a referentially opaque position, but not so in ‘I heard Pat sing’; see Higginbotham (1983), Vlach (1983). So one wants a theory that explains *why* the differences are significant in the ways that they are. And it’s not clear that one can give deflationary descriptions of how the meaning of a sentence with *covert* constituents (that are somehow linked to overt constituents) is compositionally determined; see Pietroski (2000), Collins (2003).

262 COMPOSITIONALITY, MEANING, AND CONTEXT

composition principles. So a theory with axioms like ‘*easy* means EASY’ or ‘*x* satisfies *easy* iff *x* is easy’, or theorems like ‘*x* satisfies *called the senator* iff *x* called the senator’, may be a poor theory of meaning. Even if such axioms/theorems are true, it is a tendentious hypothesis that they are formulated in the right way for purposes of explaining how humans understand language.

2.3 Radical Interpreters would Misinterpret

One might claim that negative facts are irrelevant to theories of meaning, because: (i) there *could* be creatures who correctly associate English signals with interpretations, while having no views about whether additional interpretations of the signals are possible; and (ii) such creatures would understand English, despite being different from human speakers. But it is hardly obvious that (i) is true in any theoretically interesting sense of ‘could’. And (ii) is blatantly question-begging. Why suppose that the right degree of abstraction for the study of (how humans understand) natural language is one that abstracts away from differences between us and the imagined creatures? Or put another way, why assume that for purposes of theorizing about meaning, it is *irrelevant* that human children do not grow up to be such creatures? Nonetheless, a well-known line of thought starts with this assumption, and then invites the conclusion that theories of meaning are theories of truth. Unsurprisingly, the reasons for rejecting this line of thought are closely related to some reasons for adopting a Chomsky-style conception of meaning.

The semantic properties of sentences are often said to be somehow constructable from, or at least determined by, facts concerning how utterances of those sentences *could be assigned interpretations* by a rational being in the position of someone learning the language. One imagines an alien trying to figure out, on the basis of limited evidence, what speakers are saying. By stipulation, the alien imposes only very general constraints (of rationality) on possible interpretations, and he appeals only to certain kinds of publicly available evidence; where this evidence is, at least “in principle”, available to those who naturally acquire the language. Given a suitably generous conception of availability—according to which children *could* consider native speaker reports, across various languages, about the (im)possibility of certain interpretations for certain signals—the resulting thesis would be a version of verificationism about linguistic meaning. But if the alien is supposed to be an idealized version of a child (as “field linguist”), appealing only to evidence of a

sort that typical children respond to in the course of language acquisition, the result is much less plausible than verificationism.⁶

The alien represents an arbitrary and unstable mix of two perspectives on linguistic meaning: that of a *scientist*, who will assume as little as possible about the space of possible interpretations, whatever they turn out to be; and that of a *child* who needs to figure out, on the basis of limited evidence, how signals are associated with interpretations in a given community. But actual scientists do not refuse to consider data unavailable to children. And poverty of stimulus arguments, based on negative facts like those discussed in section 2.1 strongly suggest that actual children impose substantive constraints on the space of possible interpretations for signals. Indeed, one way of summarizing the conclusion of such arguments is that human children are *not* alien interpreters of the imagined sort; see Chomsky (1969, 1993, 1995a, 2000a).

So the alien is not a reasonable idealization of *any* reasonable creature trying to figure out what expressions mean: not children, who impose arational constraints on interpretations; and not linguists, who do not impose alien limits on evidence (but rather try to figure out, by doing ordinary science, what constraints humans impose on interpretations). Nor does the alien try to determine *what portion* of the publicly available evidence actual children exploit. Instead, he considers *all* the evidence of some preferred kind. The alien thus embodies the theoretical assumption that any heuristics used by children are in principle dispensable, in favor of a more intensive search of the preferred evidence. But this assumption is quite implausible. If *all* normal children in each linguistic community converge (despite varied experience) on agreement about what signals cannot mean, this suggests that such agreement is due to properties of children, not properties of their environment. Correlatively, the challenge presented by negative facts is not “merely” to explain how *some learner could* acquire English on the basis of evidence available to a hyperattentive child. The deeper challenge is to explain how *all Englished children do* acquire English, despite variability in experience and attentiveness—and likewise for every other human language.

If semantic properties of expressions are what the alien takes them to be, this raises the question of how actual humans manage to understand human

⁶ Thinking of the child as a “little linguist” can thus be misleading. Quine (1960) restricts attention to a speaker’s disposition to endorse or reject sentential utterances. Davidson (1984) is less behavioristic, allowing for talk of speakers “holding sentences true” in a given situation (and wanting other sentences to be true). And there are alternatives to Davidson’s “principle of charity”; see e.g. Grandy (1973). But these differences of detail do not matter if the underlying idea of “meaning as (radical/alien) interpretability” is fundamentally misguided.

264 COMPOSITIONALITY, MEANING, AND CONTEXT

languages. Some philosophers have been led to the extraordinary conclusion that we don't, at least not fully; see Dummett (1986). One can hypothesize that English is a "Communal-Language" that each speaker of English imperfectly grasps. But then one needs to argue that theories of understanding are primarily concerned with Communal-Languages, and not idiolects that speakers do understand, despite Chomsky's (1986, 1996, 2000a) reasons for not appealing to Communal-Languages (or what he calls "E-languages") in explaining how speakers understand natural language; cf. Burge (1989).

Moreover, whatever speakers understand, there is no independent reason for deferring to an alien conception of understanding. *Prima facie*, the alien would not conclude that (7) has no reading on which it implies that the millionaire is from Texas,

(7) The millionaire called the senator from Texas

(I assume that the alien, like the child, does not ask which strings of words are unambiguous in which ways.) But we theorists should not conclude that (7) may have a millionaire-from-Texas reading after all—or that there is no negative fact about (7) that a semantic theory should explain, unless the alien *would* conclude that (7) has no such reading.

The alien, recall, is an epistemic monster (half-child, half-scientist) charged with a bizarre task: given a restricted body of evidence, and minimal constraints on interpretation, find the class of viable interpretations for certain noises. This imagined exercise may have some interest, since it may help reveal the size of the *gap* between typical human experience and actual linguistic competence. But for just this reason, we should not say that the semantic facts about natural languages are determined by what the imagined alien would conclude. This hypothesis about the nature of semantic properties is evidently false, given that humans cannot understand complex expressions as having certain perfectly coherent readings. One might be led to the opposite conclusion by a chain of reasoning like the following: theories of meaning are theories of truth; so we need a conception of meaning/understanding according to which theories of truth can be theories of understanding; and the best such conception treats meaning as alien interpretability. But if the last premise is correct, we have an argument against the first.⁷

⁷ One can say that a *philosophical* thesis about meaning is immune from such criticism. I don't see how any such (alleged) thesis could be evaluated in any nonstipulative way; see Chomsky (2000a) on methodological dualism. But in any case, one would have to *argue* that the thesis is relevant to the study of how human beings understand natural language. Likewise, for any "remark" about meaning offered to unconfuse a confused philosopher.

2.4 An Environment is No Substitute for a Mind

There are, of course, reasonable premises in the neighborhood. Facts about the meanings of expressions in a natural language are such that children can figure them out given evidence typically available to them. But this does not imply, or even suggest, that *any possible* learner can figure out what human sentences mean given *only* the evidence that is available to *actual* learners. In describing semantic facts, theorists will abstract away from many individual differences; and presumably, the semantic properties of an expression are not (in principle) detectable by only one person. (Even if each speaker has her own idiolect, this does not preclude understanding; two speakers can “share” a language, in the sense of associating signals with interpretations in the same way, modulo differences that are irrelevant for certain purposes.) But this does not imply, or even suggest, that the theoretically best level of abstraction will be public in the sense of ignoring any individual similarities not discernible by the alien.

If humans share a biology/psychology that imposes substantive constraints on how linguistic signals can be associated with interpretations, an inquirer will not discover this simply by attending to evidence available to any child without precocious investigation. Correlatively, the alien will be forced to look *in the wrong place* for semantic regularities that are due to constraints imposed by human biology/psychology. In his search for sources of intersubjective linguistic stability, especially with regard to how speakers use expressions to make true claims, the alien will be forced to exaggerate the relevance of the fact that speakers inhabit a shared environment. Put another way, the alien will be led to blame the environment for (i) certain aspects of intersubjective linguistic stability that are indeed due to the environment, *and* (ii) certain aspects of intersubjective linguistic stability that are due to aspects of human nature rendered invisible by alien restrictions on what is potentially relevant to theories of meaning.

This will lead the alien, who knows model theory, to favor interpretations of human speech that associate names with hunks of the environment that can satisfy predicates. Facts of the sort discussed in sections 3 and 4, which tell against such interpretations, will be ignored entirely or discounted as “noise”; for the alien has no conceptual room for the possibility that such facts reflect relevant but invisible aspects of human nature. So he may well conclude, given the constraints imposed by *his* task, that the best interpretations are those that take the form of Tarski-style truth-theories. But none of this shows that semantic theories for natural languages should associate words like

266 COMPOSITIONALITY, MEANING, AND CONTEXT

‘water’ and ‘France’ with things in the environment. Expressions may have objective interpretations that are not—and do not involve relations to—Fregean *Bedeutungen*. I return to this point.

This is not to deny that talk of “triangulation” has its place; see Davidson (1989, 2001). The environment is surely responsible for some aspects of intersubjective stability with regard to how we *use* linguistic expressions to talk about things; although this kind of stability may well presuppose meaning. Perhaps serious investigation will eventually reveal that our shared understanding of the word ‘water’ has more to do with H₂O than with internalistic properties of the word. I doubt it. But in any case, the *alien’s* tendency to look for associations between words and “stuff referred to” is not itself a reason for thinking that a theory of meaning/understanding should make such associations. We have independent reason for thinking that the best semantic theory will be one that the alien could not endorse. And in the end, I think the facts warrant a stronger conclusion: for purposes of theorizing about meaning/understanding, the best degree of abstraction will be one that (a) de-emphasizes facts about what speakers actually refer to when using language and (b) highlights internalistic properties of expressions that can be used in different environments to make semantically identical but truth-evaluably distinct claims. If this is correct, we must re-evaluate the influential idea that linguistic meanings should be specified in terms of context-sensitive rules for determining truth/reference/satisfaction conditions.

3 Systematicity with Flexibility

As Davidson (1967*a*, 1967*b*, 1984) noted, there seem to be counterexamples to his hypothesis that there are Tarski-style theories of truth for natural languages. My own view is that the constructions that Davidson himself addressed—action sentences with adverbial modification, and discourse reports—present no special difficulties, given subsequent developments of his seminal proposals.⁸ But the deeper worry has less to do with specific constructions, and more to do with some ubiquitous features of natural language, illustrated with sentences like (1).

- (1) France is hexagonal, and it is a republic

⁸ Focusing on adverbial constructions turned out to be extremely productive. And the basic framework has many virtues. See Larson and Segal (1995) for a wide-ranging and theoretically interesting illustration of a Davidsonian semantic program.

These examples, which also reveal subtle and interesting interactions between lexicalization and compositional effects, invite a Chomsky-style view. For as we shall see, this lets us retain many virtues of Davidson's conjecture without losing descriptive adequacy; cf. Austin (1962).

3.1 Lexical Flexibility: Implication and Typology

Utterances of (1) can be true. Imagine that the speaker is objecting to a crazy view according to which the shape of a country explains the local form of government. If there is a theory of truth for English, one might expect it to have a theorem according to which an utterance of (1) is true iff something in the relevant domain satisfies the following conditions, at least relative to the context of utterance: it is France, it is hexagonal, and it is a republic. But one might also suspect that nothing is both hexagonal and a republic. For even if one grants that the terrain of France is hexagonal, one might deny that the French terrain is the French republic; and one might think that republics, whatever they are, cannot be hexagonal. Moreover, even if there are such things, one might be suspicious of a *semantic* theory according to which it *follows* from (1) that there are hexagonal republics. Such a theory apparently *mischaracterizes* the meaning of (1). For a competent speaker who asserts (1) might deny that at least one thing is both hexagonal and a republic. And a competent interpreter can deny that there is any such thing, while still taking the speaker's assertion to be true.

Chomsky (1977, 2000a) offers many examples of this sort. There are, of course, possible responses. Speakers may explicitly deny what they tacitly assume; and true theories of truth may not have the implications that critics expect. But the form of Chomsky's argument is familiar. One observes that if a certain kind of semantic theory is correct, sentences would have elementary implications that competent speakers evidently do not recognize. And at least *prima facie*, this tells against theories of that kind. Consider an analogy.

According to many philosophers and some linguists, a theory of meaning for English should not imply that an utterance of 'Some bottles are red' is true only if: *there is a set* whose elements are all and only the bottles; or *there is a property* of redness instantiated by all and only the red things. Any such theory seems to mischaracterize the semantic properties of the mundane sentences. This may be harmless for many purposes, and useful for others. But intuitively, the *meaning* of 'Some bottles are red' does not ensure that an utterance of this sentence is true only if the world contains something—a set, property, or whatever—that is intimately related to but distinct from each

268 COMPOSITIONALITY, MEANING, AND CONTEXT

bottle. One can illustrate this point vividly with examples like ‘Some sets are nonselfelemental’; though a theory according to which this true sentence has provably false implications (e.g. that there is the *set* of nonselfelemental things) is just a special case of a theory that mischaracterizes meanings. There are potential responses. But these must be evaluated on their merits, not as resolutions of a paradox that any semantic theory will face. For we need not take it as given that a theory of meaning will associate ‘bottle’ with an abstract entity distinct from each bottle (and likewise for every other predicate, including ‘set’); see Boolos (1998).

Returning to (1), speakers do recognize that it implies (1*a*) and (1*b*).

(1*a*) France is hexagonal

(1*b*) France is a republic

So we want to know why the inference from (1) to (1*c*) is not trivial in the same way.

(1*c*) France is a hexagonal republic

The point is not that (1*c*) is analytically false or meaningless. Natural languages are not like formal languages with formation rules any violation of which results in gibberish; see Higginbotham (1985). Indeed, an utterance of (1*c*) in the right context might be true. Still, (1*c*) is weird in a way that calls for explanation, and likewise for ‘Something is a republic that is hexagonal’. An obvious thought is that by virtue of their meanings, ‘hexagonal’ and ‘republic’ cannot be comfortably combined to form a complex (presumably conjunctive) predicate; compare ‘green idea’ and ‘sleep furiously’ in Chomsky’s well-known example. This isn’t yet an explanation. But we can at least encode the explanandum by saying that the two monadic predicates, which are alike in some semantically relevant respects, are associated with variables of different types indicating different kinds of linguistic features. Then we can try to provide theories of how such features can be combined to create complex expressions that can be used (in ways natural for humans) to make various kinds of claims.

By contrast, it’s not clear how to even start describing the relevant facts given a theory according to which ‘France’ is semantically associated with a language-independent entity that can satisfy both ‘hexagonal’ and ‘republic’. On such a theory, ‘hexagonal republic’ is on a par with ‘brown dog’. For certain elementary purposes, we can say that all of these expressions are predicates of type $\langle x, t \rangle$, thus signifying a certain relation to names (for things in some domain) and truth-evaluable expressions. But it hardly follows

that this exhausts the semantic typology of monadic predicates. And as noted in section 2, associating ‘France’ with something referred to by using ‘France’ may be the wrong kind of theoretical abstraction. If we assume that some hunk of the environment is the source of (stable intersubjective) semantic properties of ‘France’, we may obscure significant distinctions that a semantic theory should highlight. So perhaps we should focus less on the things we use ‘France’ to talk about, and more on whatever properties of ‘France’ make it possible for us to use a name of this sort in the ways we do use such names.

3.2 Flexible Meanings and Complex Concepts

The right conception of linguistic meaning may still be undreamt of. But even if one agrees that ‘*France*’ denotes *France* is not a theory, it is at least a gesture in the direction of a theory that associates expressions with things referred to. So if only to loosen the hold of this idea, it may be useful to wave hands in an alternative direction: perhaps the meaning of an expression is an instruction for creating a concept from available mental resources.

To be sure, mentalistic conceptions of meaning are often combined with objectionable claims: that meanings are ideas, in a way that precludes successful communication; or that most words can be defined in terms of a relatively small stock of basic concepts, with the result that there is much more analyticity than there seems to be; or that there are no distinctively linguistic constraints on how humans associate signals with meanings.⁹ But these claims need not be part of a Chomsky-style view. A related point is that many theorists, including some in the “generative semantics” tradition, have explored the idea meanings *are* mental representations (or “conceptual structures”) somehow associated with expressions of natural language; see, for

⁹ See Baker (1988, 1997) for a paradigm illustration of how insights from generative grammar can be incorporated into a view according to which syntax imposes severe constraints on interpretations, thereby explaining many negative facts of the sort discussed in s. 2. See Pietroski (2003*b*) for a discussion of analyticity in these terms. While I sympathize with Jackendoff (1990) in some respects, he stresses negative facts less than I do; and I think he underestimates the role of autonomous linguistic constraints on how expressions can(not) be associated with “conceptual structures”. As he notes (p. 32), it has long been clear that features of expressions are “inadequate to the full task of conceptual description”. But theories of meaning need not be theories of concepts any more than they need be theories of truth (or referents). And since Jackendoff (1997) regards ‘Fritz is a cat, so Fritz is an animal’ as a paradigmatic case of valid inference, I assume that he is interested in (alleged) conceptual relations that go well beyond formal validity and relations established by expressions themselves.

270 COMPOSITIONALITY, MEANING, AND CONTEXT

example, Katz and Fodor (1963), Lakoff (1970, 1987), Jackendoff (1990, 1997, 2002). But my suggestion is not that linguistic expressions have *Bedeutungen* that are mental as opposed to environmental.

Recall that Strawson (1950) urged us to characterize the meaning of a referential device R in terms of “general directions” for using R on particular occasions “to refer to or mention particular objects or persons,” *and not* in terms of some entity allegedly denoted by R. This leaves room for talk of concepts as well. For we can try to characterize the meaning of R in terms of general directions for using R (on particular occasions) to express particular concepts and refer to, or mention, or think about specific things. Correlatively, one can avoid familiar pitfalls while adopting a view like the following: the word ‘France’ has certain features that get correlated in human minds with certain conceptual capacities, like the capacities to think about spatio-temporal coordinates, and about intentional properties of people who create institutions; these may not be capacities to think about (properties of) the *same* mind-independent things; and these capacities may themselves be complex and varied, in ways that tell against the claim that ‘France’ either denotes something or ambiguously denotes some things.¹⁰

Speakers can use ‘France’ to refer to various things—certain terrain, a particular nation, or whatever. Correlatively, predicates like ‘hexagonal’ and ‘republic’ seem to differ in kind. This does not yet show that *semantic* theories should mark such distinctions. Perhaps we should diagnose such facts as

¹⁰ As Lewis (1972) noted, theories that “merely” associate expressions with (instructions for creating) mental representations do not associate sentences with *truth* or *truth-in-a-model*, in the way familiar from Frege–Tarski–Montague treatments of formal languages. This was taken to be a defect of such theories. But in retrospect, one might take it to be a virtue; and as Lepore (1983) argued, appeals to truth-in-a-model provided less than advertised; see also Higginbotham (1990). The trend towards externalism, invited by alien-interpretability conceptions of linguistic meaning, seemed to support rejection of mentalistic/internalistic conceptions. But as discussed in section four, a Chomsky-style view is not threatened by Twin-Earth thought-experiments, and it is fully compatible with Kripke’s (1980) insights. Prior to “quantifier-raising” conceptions of grammatical structure, and appeals to “LF” as a level of natural language *syntax*—see May (1985), Higginbotham and May (1981), Chomsky (1981, 1995*b*)—it was also assumed that (as Frege–Russell–Wittgenstein had argued) grammatical form diverges significantly from logical form, even for relatively simple quantificational constructions like ‘The dog saw every cat’. And this suggested, wrongly, that meaning was importantly independent of natural language syntax; see Pietroski (2003*c*) for review and discussion (see also Neale, 1990, 1993; Hornstein, 1984). So if the evidence now suggests that theories of meaning/understanding are *not* plausibly viewed as theories of truth, perhaps we should conclude that a reasonable conception of meaning was abandoned for a cluster of (what turned out to be) bad reasons.

reflections of what speakers know about France, and not what they know about ‘France’; see Fodor and Lepore (1992, 2002). But we should also consider the following possibility: lexicalization is a process in which diverse mental representations can be linked via the language system. Perhaps without lexicalization, representations that are different in kind cannot be combined to form a complex concept that is usable in human thought, but (luckily for us) the language system provides resources for creating certain “common denominators”, which make it possible to create endlessly many complex mental representations with constituents that are typologically disparate. This leaves ample room for grammatical expressions that *don’t* provide ways of forming complex mental representations that are usable in fully natural ways by human minds; though such expressions may still trigger “degraded” mental representations that can be used in limited ways. The familiar idea would be that words often “fit together” in ways that the corresponding concepts by themselves do not; and in endlessly many though not all cases, linguistic expressions can be given natural interpretations.¹¹

As a first approximation, one might think about a complex monadic predicate like ‘brown dog’ as (*inter alia*) an instruction for creating a monadic concept from disparate mental resources. We may naturally think about colors as properties *of* surfaces, and think about dogs as things that *have* surfaces, with the result that brute concatenation of the relevant concepts would be unnatural for us. And there is independent reason for thinking that natural language provides a constrained system of grammatical features that can be used as rough indicators of various possession relations; see, for example, Uriagereka (2002). Perhaps such features serve as “adaptors” that make it possible for us to connect concepts of different types, thereby forming the kinds of complex concepts that we regularly deploy in ordinary human thought.¹²

¹¹ And so not a “category mistake” (cf. Ryle, 1949) if this implies that the resulting expression is nonsense or contradictory; see also Evans’s (1982) talk of “generality constraint” on distinctively human thought. Carruthers (2002) argues—drawing on lots of evidence from psychology, especially Hermer-Vasquez *et al.* (1999), Spelke (forthcoming)—that the language system plays something like this role in cognition. Variations on this theme underly a great deal of work in linguistics, both in lexical semantics (see e.g. Jackendoff, 1990; Pustejovsky, 1995; Levin and Rappaport-Hovav, 1995; Bloom, 2000) and appeals to “type adjustments” in compositional semantics (see e.g. Montague, 1974; Partee and Rooth, 1983). One need not agree with the details of such work to think that it is getting at *something* important about how the human language system relates to other human cognitive systems.

¹² Independently, it seems that some grammatical features of expressions (markers for case, person, number, etc.) do not reflect the basic architecture of the recursive system that allows for

272 COMPOSITIONALITY, MEANING, AND CONTEXT

If the idea of predicates as instructions for creating concepts seems foreign, it is worth noting—as Chomsky (1996, 2000a) does—that one can also think of ‘brown dog’ as an instruction (to the human articulatory system) for creating a complex sound. Indeed, an expression of a spoken language may just *be* a pair of instructions for creating a sound of a certain sort and a concept of a certain sort. Sentences can be viewed as instructions for pairing sentential sounds with sentential concepts.¹³ And just as phonologists can try to explain relations of rhyme in terms of relations between certain instructions for creating sounds, semanticists can try to explain relations of entailment (say, between ‘Fido is a brown dog’ and ‘Fido is a dog’) in terms of relations between instructions for creating sentential concepts. One can use the apparatus of model-theory to characterize such relations. But the utility of this apparatus is not an argument for the *hypothesis* that entailment is best explained in terms of truth (or truth-in-a-model). We may pretheoretically characterize many semantic relations in terms of truth. But often, good explanation requires redescription of explananda in overtly theoretical terms. And notions like ‘truth’ may not make the right theoretical cuts for purposes of explaining the facts that semantic theories explain; see Hornstein (1984) for related discussion.

composition of expressions. So an obvious thought, developed by Chomsky (1995b, 2000b) and others, is that such features reflect modifications of a simpler (and perhaps nontransformational) system that became more usable as a device that “interfaces” with other cognitive systems. If this is plausible, it invites a more general conjecture: apparent quirks of the human language system—aspects of the system not required in order to recursively associate signals with interpretations—reflect a natural history in which a “minimal” system has been supplemented with devices that allow for the creation of expressions with cognitively useful properties; where such expressions are interpretable as instructions for creating complex concepts that are otherwise unavailable for natural use. See Hauser *et al.* (2002), Uriagereka and Piatelli-Palmarini (forthcoming).

¹³ The details depend on the logical forms associated with sentences of natural language. But at least many sentences (like ‘Brutus stabbed Caesar with a knife’) can be viewed as instances of the logical form ‘ $\exists x\Phi x$ ’: existential closure of a monadic predicate. Any such sentence can be treated as an instruction for creating a complex monadic concept C, and then the corresponding thought of the form ‘ $\exists x\Phi x$ ’ with C as its main constituent. Elsewhere, I have argued that this paradigm covers far more than one might have thought, including quantificational examples like ‘No theory covers every case’; see Pietroski (2002, 2003c, 2004). Indeed, my suspicion is that natural language is fundamentally a system that allows for combination of monadic predicates (and a small number of relational notions associated with “thematic roles” that are in turn associated with certain grammatical relations); see also Castañeda (1967), Parsons (1990), Schein (1993, 2002, forthcoming), and Baker (1997).

From this perspective, ‘dog’ provides instructions for accessing one or more concepts already available for natural use. This leaves open various issues about the concepts: are they structured or atomic; how many ‘dog’-concepts do we have; is the species-concept dependent on the individual-concept (or vice versa, or neither); etc. Similarly, one can say that a name like ‘France’ provides instructions for accessing one or more singular concepts, which can be used to think about the various things that can count as France. This remains a vast oversimplification. But the idea, offered here as a quick illustration of *an* alternative to the idea that ‘France’ denotes a hexagonal republic, would be as follows: a speaker using (1a) can use ‘France’ to indicate a concept with which a human can think about something that has geometric properties; a speaker using (1b) can use the same lexical item to indicate a concept with which a human can think about something that has political (and hence, intentionally characterized) properties; and ‘France’ itself is an expression that makes a certain range of singular concepts available for use in the construction of various sentential concepts that can be used to make various truth-evaluable claims. But ‘France’ is not *ambiguous* in the way that ‘bear’ is. And in so far as there are French things that speakers cannot naturally refer to by using ‘France’, there are negative facts to explain. So the point is not merely that the relation of word-sounds to concepts is one-to-many. The facts are more subtle, interesting, and potentially revelatory of human thought.

In thinking about examples like (1), it is worth remembering that Plato’s poverty of stimulus argument in the *Meno* involved geometry, and that humans understand words like ‘triangle’ in a very interesting way. We know that perceptible figures can count as triangles; and we can perceptibly distinguish triangles from circles and squares, which can be drawn in the sand (say, for purposes of illustrating a generalization). Yet we also know that “real” triangles, described by theorems of geometry, are imperceptible. Hypotheses about natural language must cohere with such facts, and the fact that utterances of ‘Triangles are perceptibly different from circles’ and ‘Triangles are imperceptible’ can be true, while utterances of ‘Imperceptible triangles differ perceptibly from imperceptible circles’ are not.

This invites the thought that linguistic meanings are involved in making it possible for humans to connect percepts with a *capacity* for abstract thought that would lie “untriggered” if not for the language faculty. Perhaps we could not think about (the various things that can count as) triangles, as opposed to merely being able to classify certain things as triangular, without two integrated and integrating capacities: an ability to lexically connect concepts

274 COMPOSITIONALITY, MEANING, AND CONTEXT

corresponding to perceptual prototypes, an abstract notion of space, and the idea of proof or necessity; and an ability to create sentential concepts unavailable without mediation by linguistic expressions that have the right features. Hermer and Spelke (1994, 1996), Hermer-Vasquez *et al.* (1999), Spelke (2002), and Carruthers (2002) provide arguments that this is not mere rationalist speculation. For example, there is evidence that prelinguistic children—and animals without a language faculty—lack the ability to create some relatively simple structured concepts whose constituents are readily available. In any case, we should be wary of semantic theories according to which (i) linguistic meaning cannot play this kind of role in human *thought* because (ii) the relation between meaning and *truth* is relatively simple. This makes the study of thought and ontology even harder than it already is.

To repeat, the view is that speakers can use ‘France’ to make a variety of true claims, and that this kind of usage is possible in part because the meaning of ‘France’ does *not* associate the name with an environmental entity; though typically, a speaker uses ‘France’ to refer to something, like a government or a sports team. This view is often associated with Wittgenstein (1953). But as Chomsky (1966, 2000a) notes, it is also what traditional rationalists should expect, absent a beneficent deity or wildly optimistic assumptions about the history of natural selection. Nativists should be unsurprised if commonsense thought and talk does not reflect the structure of the world, except perhaps by sheer luck; see McGilvray (1999). Even when speakers/thinkers use language in consciously regimented ways, with the express aim of trying to describe the world (as opposed to engaging in ordinary talk about the passing show), success is not guaranteed. On the contrary, success seems to be possible only in certain domains; see also McGinn (1993). Indeed, given poverty of stimulus arguments, why should *anyone*—apart from alien interpreters and naive empiricists—expect the structure of reality to “fit” the structure of our natural ways of talking/thinking about it? Chomsky offers a more realistic conception of language, without describing the intricately structured and highly constrained phenomena of understanding with unhelpful slogans like “meaning is use”.

3.3 Typology and Ontology

We can speak quasi-commonsensibly of France being a “truth-maker” for (1), (1a), (1b), and (1c). But it does not follow that France is an entity that satisfies the predicates in these sentences, at least not if ‘satisfies’ is used in its standard technical sense, derived from Frege (1879, 1892) and Tarski (1933).

This does not challenge the commonsense claim that France is a perfectly real country that one can visit. But ordinary vocabulary is often ill-suited to the task of describing the world in theoretically perspicuous ways. Correlatively, given some technical terminology that was introduced to talk about languages *invented* for purposes of describing the world in theoretically perspicuous ways, there is no guarantee that this technical terminology will also provide theoretically perspicuous ways of describing languages with ordinary vocabulary.

One can hypothesize that France is an “all-purpose thing”, which somehow incorporates all the potential truth-makers for claims of the form ‘France is Φ ’, as part of a proposal about the *semantic* properties of ‘France’; cf. Meinong (1904) on squarable circles. But there are limits on what one can plausibly posit given the available evidence. And one can achieve implausibility without positing subsistent but nonexistent squarable circles. Put another way, if one sees nothing wrong with theories of meaning that posit Fregean *Bedeutungen* with both perimeters and politicians, one needs to say what (if anything) is wrong with more overtly Meinongian semantics. With this background in mind, I think comments like the following seem quite plausible, and in no way a denial of commonsense realism.

As far as is known, it is no more reasonable to seek some thing-in-the-world picked out by the word ‘river’ or ‘tree’ or ‘water’ or ‘Boston’ than to seek some collection of motions of molecules that is picked out by the first syllable or final consonant of the word ‘Boston’. With sufficient heroism, one could defend such theses, but they seem to make no sense at all. Each such usage of the words may well pick out, in some sense, specific motions of molecules and things-in-the-world (the world as it is, or is conceived to be); but that is a different and entirely irrelevant matter. (Chomsky, 1996: 48.)

One *could* provide a “model” of English that associates syllables with *Bedeutungen*, treating words compositionally. So if one agrees that such a model would not teach us much about linguistic meaning, the question is how much more we learn from standard models.

Similar remarks apply to pronouns. If ‘it’ has ‘France’ as its antecedent in (1),

- (1) France is hexagonal, and it is a republic

and ‘France’ does not have a *Bedeutung*, neither does ‘it’. But there are independent reasons for thinking that antecedence is a *grammatical* relation; associating a pronoun with the same entity as another expression is neither

276 COMPOSITIONALITY, MEANING, AND CONTEXT

necessary nor sufficient for antecedence. (Consider ‘My square circle has a perimeter equal to its diameter’ and ‘He must be Bob, since he is driving Bob’s car’; see Higginbotham, 1983, 1985.) So with regard to (1), an obvious thought is that while a pronoun makes the full range of its antecedent’s features available for predication—modulo restrictions, like gender or number, imposed by the pronoun—only *some* of these features will be semantically “activated” by any given predicate. In (1), ‘republic’ may indirectly (i.e. at the occurrence of ‘it’) activate features of ‘France’ that cannot be naturally combined with ‘hexagonal’. Examples like (15) provide further illustrations.

- (15) The red book is too heavy, although it was favorably reviewed, and the blue one is boring, although everyone is reading it

A speaker can utter (15), talking about which book to bring on a trip, and say something true. But it does not follow that one satisfier of ‘book’ satisfies ‘red’, ‘too heavy’, and ‘favorably reviewed’, while another one satisfies ‘blue’, ‘boring’, and ‘everyone is reading it’; see Chomsky (2000a), Pietroski (2003a). There are books; some are red, and some are heavy. But this does not imply that ‘book’, a predicate of *natural* language, has *satisfiers* at all—much less that books can satisfy other predicates that show signs of being typologically disparate. Prima facie, this theoretical claim mischaracterizes how speakers understand ordinary discourse. To be sure, there are important differences between predicates and their arguments. But this claim is detachable from more tendentious claims about how meaning is related to truth. Likewise, we can say that predicates are apt for use as devices for classifying, while names are apt for use as devices for referring, without saying that names denote satisfiers of predicates.

3.4 Ontology and Meaning

Given examples like (2) and (3), it seems that desperate measures will be needed to maintain the ontology required by theories of truth that can serve as theories of meaning for natural languages.

- (2) The government does little for the sake of the average American, whose children will inherit the massive deficit that is accumulating.
 (3) Hamlet lived with his parents in Denmark.

Other things being equal, one expects a (nontrivial) theory of truth for English to imply that an utterance of (2) is true only if the world includes: something that is a massive deficit, accumulating, and inheritable; something

that has a sake, is the average American, and has children; and a government that does little for the sake in question. So one might suspect that no such theory will be true. And one can suspect this, while conceding that some paraphrases of (2), like (2*a*), don't raise all the same concerns.

- (2*a*) The members of the current administration do little for average Americans, whose children will inherit a massive deficit due to current policies.

A theory of truth/meaning for English must associate (2), which is itself a perfectly good sentence of English, with a truth-specification. And one does not provide such a theory simply by saying that, for purposes of assigning a truth-specification, (2) is somehow associated with a more ontologically respectable paraphrase. One has to say *how* a speaker who understands (2) associates this sentence with the preferred paraphrase (and not others).¹⁴

It is often said that an utterance of (3)—sincerely produced by someone who knows that Hamlet is a fictional character in Shakespeare's play—is true iff *in the relevant story*, Hamlet lived with his parents in Denmark. Such biconditionals may well be true, at least as idealizations, and they may explain *something*. Certainly, utterances of (3*a*) can be true.

- (3*a*) In Shakespeare's famous play, Hamlet lived with his parents in Denmark.

But *prima facie*, for any sentence S, the meaning of 'In the relevant story, S' depends on the meaning of S. So I don't think appeals to 'In the story' operators will help provide a theory of truth/meaning that accommodates (3), unless one adopts Lewis's (1986) view according to which both meanings and stories are characterized in terms of Lewisian possible worlds: totalities of things as real as you and me, just not things that exist in *this* world/spacetime. On this view, there really *are* worlds with a fleshy Hamlet and mortal Polonius, and there really *are* worlds at which Hesperus is not Phosphorus, *pace* Kripke (1980). But if one rejects Lewis's conception of reality and reference, while still holding out for a theory of truth that accommodates (3), trouble awaits. For 'Hamlet' is either satisfied by nothing or satisfied by something.

Like many others, I think that Lewis adopts ontologically desperate measures. While his picture is coherent and interesting, in ways that Meinong's

¹⁴ Cf. Montague's (1974) treatment of quantificational constructions. Higginbotham (1985) notes the similarity of 'The average American has 2.4 children' to 'On the average, an American has 2.4 children'. But it is not clear how to extend this analogy to examples like (2).

278 COMPOSITIONALITY, MEANING, AND CONTEXT

(1904) was not, it is still incredible. That said, Lewis may have been right about what theories of truth for natural language ultimately require, across the wide range of cases he discusses. And *if* we assume that such theories are needed, no matter how implausible their ontological implications, it becomes very hard to offer principled reasons for resisting Lewisian conclusions. One is free to speculate that a correct theory of truth for English will require hexagonal republics though not Lewisian possible worlds. But then one has to say why the considerations Lewis presses (simplicity, scope, etc.) don't tell in favor of his theory, which really ends up being a theory of a *regimented variant* of English—that is, a language for which a theory of truth can be given, assuming enough ontology.

I suspect that some philosophers want it both ways: a theory of meaning/truth according to which meanings themselves relate expressions to the things that speakers use expressions to talk about, so that *understanding* an expression is already a way of being “in contact with” the world that makes our claims true; and a theory of meaning/truth without substantive ontological commitments, so that understanding an expression does not require a *theory* of that which makes our claims true. I won't try to argue here that this is a shell game. But I do want to register respect for Lewis's honesty, which led him to work out in detail what a theory of truth might actually require.¹⁵

If we set aside Lewis's view, it is *very* hard to see how a theory of truth/meaning for English could avoid mischaracterizing the meaning 'Hamlet' or 'Denmark' or both. This is, I think, one thing established by the vast literature on “fictional names”—names introduced for purposes of creating fiction. Such names exhibit the hexagonal republic phenomenon with a vengeance. We can say, truly, that Hamlet is a fictional character, and that Hamlet is a prince who at one point hallucinates and merely seems to see a dagger, but at another point (unintentionally) kills Polonius with a real sword; although the status of ghosts and witches in Shakespeare's plays is less clear, as illustrated by debates concerning Macbeth's interactions with the weird sisters. Theories of natural language must allow for the previous sentence.

Kripke (n.d.) provides an insightful starting point for a lexical semantics of fictional names. And of course, if 'France' can be used to refer to different

¹⁵ Jackendoff (1990) is laudably clear about the difference between his psychological conception of meaning and Lewis-style formal semantics. Linguists who claim to be pursuing the latter, but without regard for metaphysics as philosophers understand it, are often less clear about what their theories are supposed to be theories of.

things, we should expect the same to be true of ‘Hamlet’. Lots of other work may also find a place in an eventual account how names without a preexisting bearer can be used to make true assertions. My point is (not to disparage the literature, but rather) to note that names like ‘Hamlet’ invite treatment in terms of the hypothesis that the meaning of a name should be specified (not by associating the name with some entity, but rather) in terms of some array of features that the name makes available for a variety of uses; where using all the features at once would be ungrammatical and incoherent. Similar remarks apply to (16), utterances of which can be true.

(16) Teddy bears are in the next aisle, and the unicorns are right here.

But utterances of ‘There are no unicorns’ can also be true.¹⁶

Once we consider the possibility that (1)–(16) illustrate related phenomena that are ubiquitous in natural language, as opposed to thinking about (1)–(16) as a hodgepodge of marginal cases to be set aside, I think it becomes clear just how bold Davidson’s conjecture was. Linguistic meanings don’t seem to be functions from contexts to truth/reference/satisfaction conditions (even setting aside the technical difficulties presented by ‘Yesterday, I said that Hesperus is Phosphorus’). In section 4.3, I return to this point. But first, I want to enter some disclaimers, and briefly compare a Chomsky-style view with some alternatives that specialists will know about.

4 Caveats and More Doubts

My claim is not that we should abandon current semantic theories, or that standard textbooks are complete bunk. It is rather that “axioms” like ‘*France* denotes *France*’ are best read, despite appearances to the contrary, as preliminary claims about intrinsic features of linguistic expressions—and that we should bear this in mind, as we revise our current theories. But the suggestion is not that each word has a unique array of grammatical features. And I am not denying that causal-historical facts, of the sort Kripke (1971, 1980) and others have discussed, bear on what speakers refer to with names.

¹⁶ Generic plurals like ‘Teddy bears’ present their own complications, as Chomsky (1977) discusses—especially in examples like ‘Unicycles have wheels’, ‘Beavers build dams’, and ‘Poems are written by fools like me’; cf. Carlson and Pelletier (1995).

4.1 Internalistic Meanings, Externalistic Truths

In many idiolects, 'Latvia' and 'Estonia' may be type-identical, modulo pronunciation. Perhaps the same is true for 'elm' and 'beech', or 'arthritis' and 'rheumatism'; see Putnam (1975), Burge (1979). But in the absence of evidence to the contrary, children seem to assume that different expressions in the same language have different meanings. So I assume that except for a few special cases, marked as such, distinct expressions are understood as semantically distinct. It doesn't follow, though, that understanding 'Latvia' is a matter of knowing that it denotes Latvia. Understanding the name may be a matter of tacitly knowing that (i) the name is distinct from other expressions of the same kind, and (ii) the name has certain features in virtue of which it can be used as a device for referring to a place, perhaps characterized in terms of intentional properties.

One might say that 'Latvia' is relevantly like 'water', and that understanding 'water' is relevantly like having seen water: one cannot be in such a state without bearing the right causal-historical relation to some H_2O . There is, however, little if any evidence in favor of this prima-facie implausible thesis. I am happy to say that understanding is importantly like perception. But one might have thought that understanding 'water' is more like seeing (or hearing) the expression 'water' than seeing water; where expressions are individuated so that creatures in H_2O -less environments could perceive and use 'water'. Since it has become common to think otherwise, let me stress: Putnam and Burge never showed that *theories of linguistic meaning* must employ a notion of expression such that my Twin-Earth duplicate and I use typologically distinct expressions. (Presumably, my twin is like me with respect to intuitions that linguists care about.) The thought-experiments suggest that some facts about *how humans use language* cannot be explained in internalistic terms; and this bears on certain philosophical projects and claims. But one needs a premise to get substantive claims about meaning from these claims about use.

Referring to water is relevantly like seeing water. It can't be done without *some* kind of contact with at least some H_2O . And for purposes of figuring out what a speaker is trying to say, Davidsonian triangulation is presumably important, even if alien interpreters exaggerate its importance. I am also inclined to agree with externalists like Burge (1979, 1989), who hold that the truth or falsity of an utterance can depend on the norms of a relevant community—and notions of rational commitment—in ways not captured by the ways in which the meanings of indexical/demonstrative expressions

track certain aspects of conversational situations; cf. Fodor (1987), Stanley (2000). Externalism about truth may well be correct in this nontrivial sense. But one can deny that meaning and understanding are tightly connected to truth and reference and rationality, we need not say that *understanding* ‘water’ is relevantly like having seen water, or that ‘water’ is, by virtue of its linguistic *meaning*, an indexical or demonstrative expression like ‘nearby’ or ‘that stuff’. Perhaps ‘water’ is, as it appears to be, a mass-noun with no part that somehow indexes H_2O on earth but not Twin-Earth; and perhaps the thought-experiments just reveal that the relation between meaning and truth is not as simple as some philosophers thought (or hoped).

If intuitions about Twin-Earth thought-experiments reflect our tacit views about truth and reference—and what speakers *commit themselves to* when they use language in certain ways—then such intuitions do not tell against a Chomsky-style internalism about linguistic meaning. For the internalist view on offer is one according to which linguistic meanings *guide and constrain without determining* truth, reference, and other (norm governed) expression–speaker–world relations. This leaves room for the claim that such relations are interestingly externalistic. So one can hardly use the thought-experiments to argue for this claim, and then use them again to argue that meaning is like truth and reference in this respect.¹⁷

We can *invent* a language in which: a predicate Φ is associated with a function from possible worlds to substances like H_2O ; and some name α is associated with a function from possible worlds to either (i) all-purpose entities like the alleged hexagonal republic of France, or (ii) functions from contexts and n-tuples of Fregean *Bedeutungen* to purpose-specific entities. This may establish the coherence of corresponding hypotheses about natural

¹⁷ One can try to provide independent arguments that the connection between meaning and truth is tighter. But the premises must be more plausible than the claim that theorists should defer to the alien with respect to human understanding. Burge (1979) and others offer arguments that rationalizing explanations of human action unavoidably traffic in externalist notions of *intentional content*; cf. Fodor (1987). I happen to find this conception of human action plausible; see Pietroski (2000). But why think that linguistic meaning is like intentional content in this respect, especially since the study of the former has delivered better theoretical results than the study of the latter? For all we know, human intentional content may itself be an interaction effect one of whose determinants is (internalistic) linguistic meaning. Of course, one would like an account of how it all hangs together—thought, communication, meaning, reference, truth, confirmation, atoms, constellations, praise, condemnation, and everything else. And one can define ‘Language’ so that a Language would have the properties needed to make it all hang together in some envisioned way. But it doesn’t follow that there are Languages, much less that they include natural spoken languages.

282 COMPOSITIONALITY, MEANING, AND CONTEXT

language; see Kaplan (1989). But plausibility is another matter. Stanley (2000, 2002) outlines an intriguing view according to which the truth of sentential utterances depends on the environment only in ways tracked by (the meanings of) overt or covert constituents of sentences. I don't think this generalization is warranted. But I won't try to argue against Stanley's view here; though see Blair (forthcoming) for an argument that the requisite covert constituents are not there, even in cases where positing them seems most plausible—for example, quantifier domain restriction (cf. Stanley and Szabo, 2000).

For present purposes, let me just say that Chomsky offers a less radical response to examples which suggest that truth depends on the environment in ways not tracked by theories of meaning for natural languages. One need not say that linguistic expressions have, in addition to all their other features, many covert indices not detectable with current tests. So Stanley's criticism of *other* responses does not yet undercut the force of all the apparent counterexamples to his very general thesis. That said, Stanley—and those he criticizes, like Bach (1994)—may be importantly right about *something*. The mental representations *indicated by* linguistic expressions may well have elements (not corresponding to elements of the sentences speakers utter) that track many ways in which the truth of utterances can depend on the environment.

There may also be symbol–world regularities not explained by theories of linguistic meaning. If there is a language of thought with its own “psychosemantics”, this is presumably relevant to questions about truth. And *perhaps*, as Fodor (1987, 1998) argues, a correct theory of meaning for Mentalese will associate primitive expressions of Mentalese with Fregean *Bedeutungen*; perhaps sentences of Mentalese even have (context-sensitive) truth-conditions, not merely truth-indications. For present purposes, I take no stand on these issues. Though for all we know, the relation between Mentalese and truth is also less than fully systematic, while Mentalese and a spoken language and communal norms *together* impose enough constraints to make truth stable and interesting (*pace* deconstructionists). Even if we don't know how, it seems clear that at least on occasions of use where we are trying to be careful, we can think and talk about the world in ways that are objectively right or wrong. But this hardly shows that any language we ordinarily use to think or talk is a language that has a truth-theory. And in any case, a theory of denotation for Mentalese would not obviate the need for theories of meaning for spoken languages, if only because of the relevant negative facts.

Fodor (1998) sometimes speaks as if one can account for all the semantic facts regarding spoken languages by saying that each sentential utterance gets associated with a token of some mental sentence; see also Schiffer (1992, 1994a, 1994b, 2000). But this doesn't begin to explain why sentences of English are *not* associated with sentences of Mentalese in certain ways. One wants to know why (7)

(7) The millionaire called the senator from Texas

cannot indicate a Mentalese sentence that is true iff the millionaire from Texas placed a call to the senator. One wants to why 'Brutus stabbed Caesar' cannot indicate a Mentalese sentence that is true iff there was a stabbing of Brutus by Caesar, and likewise for all the other facts regarding nonambiguity in natural language. These facts call for substantive assumptions about *how* grammatical structures of spoken languages can(not) be related to interpretations; see Higginbotham (1994), Matthews (2003). Since Fodor knows about negative facts, perhaps when he says that spoken languages do not have a compositional semantics, he just means that there are no systematic theories of truth for such languages.

One can define 'semantics' so that language L has a semantics only if there is a theory of truth for L. But one cannot stipulate that all the relevant explananda are explained by syntactic structures for spoken languages, a denotational semantics for Mentalese, and a mechanism that associates each sentential utterance *u* with a token of some mental sentence that expresses the thought expressed with *u*. Moreover, even if we identify the interpretations of certain "labels" in Mentalese with certain things that speakers can refer to by using certain names of a spoken language, it does not follow that the things referred to are the interpretations of spoken names.

Kripke (1980) noted that a speaker of English might see no significant difference between (i) the distinction between 'Feynman' and 'Gell-Mann', and (ii) the distinction between 'Cicero' and 'Tully'. For a speaker might have no way of distinguishing Feynman from Gell-Mann, except by recourse to metalinguistic predicates like *was called 'Feynman'*; and such a speaker might think, mistakenly, that Cicero and Tully were distinct Romans. Kripke made this observation in the context of arguing against theories according to which 'Feynman' and 'Gell-Mann' would be synonymous for such a speaker. But it also suggests that the semantic difference between these names has nothing to do with the difference between the physicists, since the semantic difference between 'Cicero' and 'Tully' cannot be even partly due to a way that Cicero

284 COMPOSITIONALITY, MEANING, AND CONTEXT

differs from Tully. And Kripke does not say that ‘Cicero’ and ‘Tully’ are synonymous.¹⁸

One can hypothesize that understanding ‘Cicero’ is relevantly like being causally related to Cicero, that names for the same thing are synonymous, and that the relation between meaning and truth is relatively simple—arguments to the contrary notwithstanding. Given such a view, many contrasts that might have been explained as semantic contrasts will have to be explained in some other way; see Braun and Saul (2002) for a proposal that engages with the difficulties, instead of just labelling them as ‘pragmatics’. But one wants to know if there is any *evidence* that motivates this conception of synonymy, given a Chomsky-style alternative, which leaves room for various projects concerned with the use of meaningful expressions (and notions of rational commitment); cf. Soames (1987, 1995). Perhaps alien interpreters would end up identifying the meanings of names with things named. But if anything, this should make us more skeptical of “direct reference” conceptions of linguistic meaning—even if we follow Frege in thinking that for certain norm-governed enterprises like scientific inquiry, we should use each expression as though its meaning is its *Bedeutung*.

Let me conclude this subsection with a brief remark about rigid designation. One need not say that ‘Aristotle’ denotes a certain long dead philosopher in order to accommodate Kripke’s insights. As a matter of causal-historical fact, speakers use this “famous-person name” to talk about a certain long dead philosopher. Speakers also tacitly know that names have both causal-historical associations and descriptive associations; that these aspects of use can conflict, with regard to “who we are talking about” when we use the name; and that in such cases, the causal-historical associations trump the descriptive associations. But the plausible hypothesis that names are devices for referring “rigidly” (and not by description) does not require the implausible hypothesis that names denote things. On the contrary, one might ask which thing ‘France’ rigidly denotes. For we can coherently describe a possible situation in which the terrain of France is not inhabited by people who have a republican form of government, and a possible situation in which the republic of France has different borders.

This perfectly familiar point again suggests that there is no all-purpose *Bedeutung* for ‘France’. Any such thing would be denoted rigidly by ‘France’;

¹⁸ On a Chomsky-style view, names may be more like predicates than “logical constants” of the predicate calculus. But this is independently plausible; see Burge (1973), Longobardi (1994).

and so, given the truth of various counterfactual claims, it would need to have *both* its geometric *and* political properties *inessentially*. But what is this alleged thing, which conveniently has all the properties something needs to have to be a truth-maker for all claims of the form ‘France is, or at least might have been, Φ ’? Was the republic of France formerly a hexagonal monarchy? Could the republic have been a communist state, or a loose confederation of anarchist associations? If not, perhaps we should detach the idea that speakers use names to perform acts of rigid reference from the idea that names have referents rigidly.

4.2 Extensionality

I do not, however, want to argue about words like ‘meaning’, ‘understanding’, and ‘semantics’. If someone insists that such words describe relations between expressions and potential objects of reference, my claim can be conditionalized: *if* such insistence is correct, there may be no theories of meaning for natural languages, since ‘understanding’ has been defined as a label for (what turns out to be) a massive interaction effect; and we do not need theories of the interaction effect to account for the facts, positive and negative, concerning how humans associated signals of a spoken language with interpretations. But whatever the terminology, we can try to provide theories of (speakers’ tacit knowledge concerning) intrinsic properties of linguistic expressions, supplemented with claims (which may not rise to the level of theories) about the use of meaningful expressions. Likewise, I don’t insist that semanticists eschew the term ‘denotes’. Theorists can and do create special contexts in which a name can be used to talk about its semantic properties. The one that has become standard—writing axioms like ‘*France*’ denotes *France*—makes it easy to ignore lexical flexibility. Such idealization is appropriate for certain purposes. And it is harmless, so long as we don’t think that invoking the term ‘denotes’ magically dispels lexical flexibility, or shows how to accommodate it in a theory of truth.

Correspondingly, I am *not* objecting to the idea that theories of meaning can be formulated in a metalanguage governed by an *extensional* logic. One way to see this point is by thinking about other ways of using names to talk about their semantic properties. We could invent a technical term ‘meanotes’ and write axioms like ‘*France*’ meanotes *France*, taking this to be shorthand for a cluster of claims like the following.

‘France’ is an expression (of a certain type) that makes certain linguistic features available for use. Speakers can use these features to perform

286 COMPOSITIONALITY, MEANING, AND CONTEXT

referential acts of various kinds (and thereby refer to things of different kinds). Given the contingent history of how ‘France’ has been used, speakers of English can use it to refer to the various things that can count as France in various contexts, as opposed to other things (like those that can count as Germany). But these contingencies may be extraneous to theories of meaning/understanding, which may turn out to be theories of (speakers’ tacit knowledge regarding) certain essential and internalistic properties of expressions. Although the contingences are relevant to questions of truth or falsity.

But then for purposes of writing down a real theory—with theorems concerning the semantic properties of complex expressions and axioms concerning the semantic properties of words—we would need a logic for claims of the form *expression* Σ *meanotes* α ; where the logic can be combined with plausible hypotheses about how meaning/understanding is related to meanoting. And providing such a logic will be hard, given the stipulations governing what *meanotes* α means in the metalanguage.¹⁹

In general, it is bad methodology to adopt a theoretical vocabulary that forces one to come up with a complete correct theory before offering any theory from which theorems can be derived. Better to let oneself write down and later revise partial theories that are false, as part of a process that might eventually lead to reasonably good idealizations that partly explain a certain range of phenomena. So we want an alternative to ‘meanotes’ that does not require a special logic. We want to offer comprehensible theories, and see where the difficulties lie, without having to worry about what follows from what. In this spirit, Davidson (1967a, 1984) proposed that, instead of trying to provide theories with axioms like ‘*France*’ *means* *France* and theorems like

¹⁹ Should the logic licence the inference from ‘*Hesperus*’ *meanotes* *Hesperus* and *Hesperus* = *Phosphorus* to ‘*Hesperus*’ *meanotes* *Phosphorus*? If so, are we saying that ‘*Hesperus*’ and ‘*Phosphorus*’ are synonymous if *Hesperus* is *Phosphorus*? Suppose one speaker uses ‘*France*’ to refer to certain terrain (knowing full well that *France* is also a republic), while another speaker introduces ‘*Gaul*’ as a device for referring to the same *terrain* but stipulating that *Gaul* is the wrong sort of thing to be a republic. Is ‘*France*’ *meanotes* *Gaul* true, false, or neither? Such questions need to be settled in order to know what a theory with axioms like ‘*France*’ *meanotes* *France* implies. Putting the point in a way friendly to Quine (1951): one can try to accommodate the facts in various ways; but opting for a theory governed by a nonstandard logic will not be one’s first choice. That said, I think there are good reasons for adopting a second-order metalanguage, which is *not* to say that the second-order variables range over sets of things that first-order variables range over; see Boolos (1998), Schein (1993, 2002), Higginbotham (1998), Pietroski (2004).

'France is a republic' means that *France is a republic*, we should try to do two things: provide theories with axioms like *'France' is true of France* and theorems like *'France is a republic' is true iff France is a republic*; and show how such theories can do the theoretical work that theories of meaning need to do.

This turned out to be a terrific methodological proposal. I fully endorse the strategy of supposing that the core semantic notions are extensional, and treating apparent counterexamples (propositional attitude reports, verbs like 'hunt' and 'worship', etc.) as special cases to be dealt with as such; see Larson *et al.* (forthcoming), cf. Montague (1974). But Davidson's replacement of 'means that' with 'is true iff' also reflected his implausible views about the nature and source of semantic phenomena; see 2.2 above. One can abandon these views and retain the practice of writing axioms like *'France' denotes France*. For engaging in the practice does not commit one to the hypothesis that a correct theory of meaning for English will associate 'France' with an entity that satisfies any predicate Φ , such that utterances of a sentence formed by combining 'France' with Φ are true. Instead, one can view the use of axioms like *'France' denotes France* as indications that certain idealizations like the following are operative:

For purposes of explaining the limited range of facts this theory purports to explain, we're ignoring a lot of what makes 'France' the expression it is—an expression that can (given contingent facts) be used to refer to the various things that can count as France. Likewise, many typological differences between predicates (including those that distinguish 'hexagonal' from 'republic') will be ignored. Indeed, all that really matters for these purposes is that 'France' is (i) a potential grammatical argument of a predicate, (ii) a word that can be used to refer to something, and (iii) semantically distinct from other words of this type, unless some other axiom says otherwise.

This leaves room for the claim that 'France' has a hexagonal republic as its *Bedeutung*. But we should be clear that this is the analog of what we would need to say, in terms of rules governing the derivation of theorems, given axioms like *'France' meanotes France*. In my view, blaming the language-independent world for the apparent gap between meaning and truth is no more plausible than blaming logic. And if axioms like *'France' denotes France* reflect idealizations that abstract away from all the reasons for thinking that there are no theories of truth for natural languages, then the use of such axioms does not even suggest that there are such theories. So one can endorse

288 COMPOSITIONALITY, MEANING, AND CONTEXT

much of the work done by theorists who use such axioms, typically as part of a scheme for encoding other more interesting claims about natural language, while remembering that the operative idealizations make it hard to use the virtues of current theories as arguments for the claim that theories of meaning are theories of truth. *If* subsequent theorizing leads to less idealized theories that are plausibly theories of truth, that will be another matter. But like Chomsky, I think the trend is in the other direction.

4.3 Lexical Flexibility and Standard-Shifting

Consider one last example of linguistic flexibility discussed by Austin (1962) and many others (see Travis, 1985, 1996). Some utterances of (1a) can be true, while others are false.

(1a) France is hexagonal

But a theory of truth for English will presumably have some theorem like the following, ignoring tense for simplicity: an utterance of (1a) is true iff the thing denoted by 'France' is Φ ; where Φ is a predicate of the metalanguage. So even setting aside worries about the alleged denotatum of 'France', there is a problem. The predicate Φ will be satisfied by whatever things it is satisfied by; and one will mischaracterize the meaning of 'hexagonal' by saying that it is a predicate satisfied by all and only those things—call them the Xs. The Xs may be the things that satisfy 'hexagonal' given *some* standards for what counts as hexagonal. But whatever the Xs are, competent speakers will know that 'hexagonal' can be used as a predicate *not* satisfied by all and only the Xs.

There are many potential replies to this kind of argument, and I cannot adequately address them here. But again, my point is more to raise the question of whether such replies are motivated, and less to argue that they are wrong. One might use 'hexagonal' itself in a theory of meaning—and not just as a temporary device to be replaced (eventually) with something else—even though 'hexagonal' is a poor candidate for a *theoretical* term, especially if France can satisfy it. But then metalanguage predicates like *satisfies 'hexagonal'* will have flexible meanings. So even if we allow for the use of such predicates in theories, despite Frege–Tarski admonitions against doing so, it seems that a “theoretical” sentence formed by combining *satisfies 'hexagonal'* with a suitable label for an entity is not a sentence that itself expresses a clear hypothesis. (This suggests that 'satisfies' is being used quasi-commonsensically, and so misleadingly.) Perhaps as used in a suitable

theoretical context, *France satisfies 'hexagonal'* is a truth-evaluable claim that can be empirically assessed as a clear hypothesis about natural language. But once a context is fixed, *satisfies 'hexagonal'* will be satisfied by whatever things it is satisfied by relative to that context; and prima facie, one will mischaracterize the meaning of 'hexagonal' by saying that it is satisfied by all and only those things.

Cappellen and Lepore (2003a, 2003b, forthcoming)—henceforth, C&L—claim that this difficulty can be avoided. On their view, theorems of the form *S is true iff p* do indeed have meanings as flexible as the meanings of the object language expressions in question; and likewise for the corresponding axioms. C&L maintain that such axioms nonetheless comprise an honest theory, which need not and should not be relativized to a context so that certain things are all and only the satisfiers of *satisfies 'hexagonal'*. I sympathize with the spirit of this proposal, which is to relieve semantic theories of the burdens imposed by the idea that theories of meaning should reflect all the ways that truth can depend on the environment. As C&L argue, one can and evidently should theorize about meaning/understanding in abstraction from many factors relevant to truth. They say, and I agree, that one should abstract away from aspects of context-sensitivity not indexed by *expressions* of the language in question. But once one accepts this point, I don't see any theoretical motivation for retaining the idea that theories of meaning are theories of truth.

My suspicion, which I won't try to argue for here, is that appeal to flexible truth-theoretic axioms amounts to vascillation between two perspectives: a Chomsky-style view, combined with a preference for encoding semantic theories in terms of *constraints on truth* imposed by expressions, as opposed to *features of expressions* that impose constraints on truth; and a much less plausible "deflationary" view, combined with the idea that a "philosophical" theory of meaning need not account for explananda that go unexplained by adopting axioms like *x satisfies 'easy' iff x is easy*. For present purposes, let me just note that appeal to *theories* whose axioms have flexible meanings is itself a nonstandard response to Wittgenstein–Austin–Chomsky examples. And if C&L offer the best alternative to a Chomsky-style view, then the initial motivations for adopting the latter are relatively clear, at least if one assumes that we do not know a priori what theories of meaning should (not) explain.

One can, of course, hypothesize that all context-sensitivity is relevantly like indexicality. Perhaps 'hexagonal' indexes standards, much as 'I' indexes speakers. But one wants to see the evidence, independent of the dogma that

290 COMPOSITIONALITY, MEANING, AND CONTEXT

theories of meaning are theories of truth. Prima facie, the flexibility of ‘hexagonal’ is importantly different than the indexicality of ‘I’.²⁰ With regard to the former, context matters because, for some things, there is no *clear-cut* answer as to whether or not they are hexagonal. But with regard to ‘I’, it is not just that there is no *clear-cut* answer to the question of whether or not some individual satisfies ‘I’ independent of how (i.e. by whom) the pronoun is being used. The question isn’t even coherent until the speaker is identified, at which point there is no question left. By contrast, as many authors have noted, standard-shifting seems more like—and may well be intimately related to—vagueness; see Graff (2000) for discussion. Though for just this reason, assimilating phenomena of standard-shifting to indexicality seems to mischaracterize both.

This is not to deny that the interesting questions about vagueness remain. One still wants to know what is wrong with the reasoning in Sorites paradoxes. But we should not assume that understanding is so tightly connected to truth that the following conditional holds: if the dependence of truth on the environment is vague and situation-sensitive, then *expressions of natural language* track even this dependence in the way that indexicals track other kinds of dependence. This assumption does not help resolve the paradoxes; and it may make them worse.

Correspondingly, even given a conversational situation, it can be vague as to whether or not a given entity is hexagonal. One can say that there are many contexts for each conversational situation, perhaps with no fact of the matter as to which is *the* context relevant to the truth of an utterance in the conversational situation. I think this gets the (one-to-many) relation between contexts and conversational situations backwards, thereby making it mysterious how contexts could be related to linguistic understanding. But in any case, with regard to borderline cases of hexagonality, there is no independent reason for thinking there was a clearer standard “there” that the speaker somehow failed to indicate; prima facie, there is nothing that would settle the question.

²⁰ Given any particular context, ‘hexagonal’ would have to index (not an entity of the usual sort, but rather) something that associates ‘hexagonal’ with some things; and one might wonder how this works, even setting aside concerns about whether it requires the paradoxical assumption that every predicate has an extension. And if the claim ends up being that contexts are (not just Kaplan-style n-tuples of potential satisfiers for indexed expressions, but also) entire possible worlds, the resulting “theory” is trivial: the meanings of sentences determine truth-conditions relative to contexts, because truth is determined by meaning and everything else relevant to truth; see Pietroski (2003*a*) for related discussion. But put these concerns aside.

Again, it's important not to be misled by the fact that we can invent formal languages governed by a supervaluationist logic. Such languages may illuminate certain aspects of vagueness. But it hardly follows that natural languages *are* languages of this sort, or that the phenomenon of natural language vagueness *is* the phenomenon of "supervaluationism"; see Williamson (1994) for trenchant criticism. So if one rejects Williamson's own conclusion—according to which 'hexagonal' and 'bald' have precise extensions, unbeknownst to competent speakers—one might conclude that Williamson (like Lewis) offers a nice *reductio* of the idea that predicates of natural language have meanings that can be correctly characterized with Tarski-style theories of truth. But this leaves room for possibility that Williamson is right about how we *ought to use* language for purposes of theorizing. It may well be that we have a "regulative ideal" according to which truth is tightly connected to the meanings of expressions in a *Begriffsschrift*. And it may well be that some surprising claims about natural language *would* be descriptively correct *if* natural language meaning/understanding was related to truth this way.

One can still maintain that the meaning of 'hexagonal' somehow determines a function from contexts to satisfaction conditions. But I don't think this is any better motivated than the idea that speakers *understand* vague monadic predicates like 'red' and 'bald' by associating such predicates with *functions from numbers to functions from entities to truth-values*; where the numbers correspond to "precisifications" of the predicate (cf. Lewis, 1972). Supervaluationist *models* of understanding may be useful for certain purposes. But as Sainsbury (1990) and others have noted, we shouldn't conclude that a word like 'bald' or 'red' is a predicate semantically associated with a function from precisifications to functions. For such a predicate is no more vague than any other predicate associated with a function. And *prima facie*, natural language predicates have flexible meanings that make it impossible to characterize their meanings in terms of functions, without ignoring their vagueness (and thus *mischaracterizing* their meanings). One can often idealize away from vagueness, but not when it comes to accounting for vagueness.²¹

²¹ See also McGee and McLaughlin (1994), Fodor and Lepore (1996), Pietroski (2003a). The general point is clear from Benacerraff (1965): if a theoretical picture forces us to say that Xs (numbers, meanings, or whatever) are things of a certain sort, Ys, but identifying any particular X with any particular Y seems to mischaracterize Xs—say, by *overdescribing* them, with consequent indeterminacy as to which Y a given X is—perhaps we should look for another theoretical picture.

292 COMPOSITIONALITY, MEANING, AND CONTEXT

That said, the caveats of this section apply. I am not saying that causal-historical facts are irrelevant with regard to what a given predicate is intuitively true of. Nor am I saying “axioms” like *x satisfies ‘hexagonal’ iff x is hexagonal* are bunk. This is one way of encoding a perfectly fine idea: ‘hexagonal’ is a monadic predicate; and given some things to talk about, such a predicate is apt for use (on a given occasion) as a device for sorting the things in a certain way, just as ‘France’ is apt for use as a device for referring (on a given occasion) to one of the things. But given some things, there are many overlapping ways of sorting them such that for each of those ways, a speaker of English can use the word ‘hexagonal’ to sort the things in that way. The question is whether we theorists should describe this fact about language use by characterizing the *meaning* of ‘hexagonal’ in terms of a mapping from things to ways of sorting them—and not in terms of *intrinsic features of the word that make it possible* to use ‘hexagonal’ as a device for sorting things in certain ways across various conversational situations. The theoretical task, as always in this domain, is to figure out *how* meaning is related to use. It is not enough to just say that each aspect of use reflects meaning; but encoding each aspect of use in claims about meaning is just a special case of ensuring descriptive adequacy at the cost of explaining nothing.

5 Concluding Remarks

Many examples tell against the idea that theories of meaning/understanding will be theories of truth. Perhaps these are all special cases requiring special treatment; one theorist’s *reductio* is another’s research program. But at some point, one has to wonder what *truth*-conditional semantics explains, over and above what can be explained without supposing that theories of meaning are theories of *truth*.²² Are there any nonspecial cases, apart from rarefied sentences like ‘Two plus three equals five’? Natural language may not fit the model of a language in which names are semantically associated with entities that are satisfiers of predicates. This was a fruitful model that allowed theorists to start accounting for a certain range of elementary facts. And it simplifies discussion, in harmless ways, when the flexibility of lexical items is not at

²² Higginbotham (1989a, 1989b) offers suggestions about what knowledge of reference might explain. But given apparent counterexamples to theories of truth, one needs to argue that such knowledge (in so far as speakers have it) is knowledge of *meaning*, as opposed to an interaction of linguistic understanding and other aspects of human cognition; see Pietroski (2003a).

issue. But the explanatory value of the model may be limited, in ways that now require attention, if we want *better* models that start to account for ways in which lexicalization and composition *interact*—and other ways in which natural languages are importantly unlike a *Begriffsschrift*.

We may have reached the stage at which the simplifications frustrate theorizing more than they promote it. One can say this while agreeing that, at an earlier stage of theorizing, it was more important to stress that the Frege–Tarski toolkit was applicable in (theoretically illuminating ways) to the study of natural language—and that natural languages are importantly like a Fregean *Begriffsschrift*. It may be convenient to express this last point, in reply to those who still deny it, by saying that there are theories of truth for natural language. But one shouldn't confuse a slogan for a plausible hypothesis. Likewise, since it is part of a theorist's job to invent hypotheses that initially seem like wild overgeneralizations, it may be convenient to remind theorists of certain facts (that really do tell against *certain* generalizations) by saying 'meaning is use'. But in fact, meaning *constrains* both use and truth, in subtle and interesting ways. An account that does justice to natural language will have to accommodate facts which suggest both that (i) use and truth are *very* complicated, perhaps in many intractable ways, and (ii) meaning is systematic and in many ways theoretically tractable, even for creatures with our limited cognitive powers.

I began this paper with a quote from Chomsky. Let me end with one from Kripke.

I find myself torn between two conflicting feelings—a 'Chomskyan' feeling that deep regularities in natural language must be discoverable by an appropriate combination of formal, empirical, and intuitive techniques, and a contrary (late) 'Wittgensteinian' feeling that many of the 'deep structures', 'logical forms', 'underlying semantics' and 'ontological commitments', etc., which philosophers have claimed to discover by such techniques are *Luftgebäude*. (1976: 412 n. 56).

Both sensibilities can also be found in Chomsky, who offers an attractive suggestion about how to resolve the apparent tension: meaning is less tightly connected to truth (and ontology and alien interpretability) than a lot of work suggests; expressions have semantic properties; but these are intrinsic properties of expressions that *constrain without determining* the truth-conditions of utterances. One can say that semantics is a species of syntax on this view. But that is not an objection. Given how form constrains meaning in natural languages, perhaps we should indeed replace the idea that semantic properties are *not* syntactic properties with a suitably expansive

294 COMPOSITIONALITY, MEANING, AND CONTEXT

view of syntax. In any case, we should take Chomsky's view seriously—instead of insisting on a conception of meaning according to which his insights can bear essentially on syntax, but only tangentially on larger questions about understanding. That is, we should make conceptual room for the possibility that natural language is unlike a *Begriffsschrift*, in that the relation between meaning and truth is looser, while the relation between meaning and form is tighter.

We should evaluate claims about linguistic meaning, truth, and context-sensitivity in light of our best theories of natural language, instead of insisting that these theories conform to externalist dogma. Truth may well depend on communicative situations in ways that should *not* be indexed by theories of meaning/understanding for natural languages. In which case, we must revise many current claims about how meaning, truth, and context are inter-related.²³

REFERENCES

- Austin, J. (1961). *Philosophical Papers* (Oxford: Oxford University Press).
 — (1962). *How to Do Things with Words* (Oxford: Oxford University Press).
 Bach, K. (1994). 'Conversational Implicature', *Mind and Language*, 9: 124–62.
 Baker, M. (1988). *Incorporation* (Chicago: University of Chicago Press).
 — (1997). 'Thematic Roles and Grammatical Categories', in L. Haegeman (ed.), *Elements of Grammar* (Dordrecht: Kluwer), 73–137.
 Barber, A., ed. (2003). *Epistemology of Language* (Oxford: Oxford University Press).
 Bäuerle, R., Schwarze, C., and von Stechow, A., eds (1983). *Meaning, Use, and Interpretation of Language* (Berlin: de Gruyter).
 Benacerraff, P. (1965). 'What Numbers Could Not Be', *Philosophical Review*, 74: 47–73.
 Blair, D. (forthcoming). *The Grammar of Context Dependence*.
 Bloom, P. (2000). *How Children Learn the Meanings of Words* (Cambridge, Mass.: MIT Press).
 Boolos, G. (1975). 'On Second-Order Logic', *Journal of Philosophy*, 72: 509–27.
 — (1998). *Logic, Logic, and Logic* (Cambridge, Mass.: Harvard University Press).
 Burge, T. (1973). 'Reference and Proper Names', *Journal of Philosophy*, 70: 425–39.

²³ For helpful comments and discussion, my thanks to: Noam Chomsky, Susan Dwyer, Norbert Hornstein, Ernie Lepore, Jim McGilvray, and Georges Rey.

- (1979). 'Individualism and the Mental', *Midwest Studies*, 4: 73–121.
- (1989). 'Wherein is Language Social', in George (1989).
- Braun, D., and J. Saul (2002). 'Simple Sentences, Substitutions, and Mistaken Evaluations', *Philosophical Studies*, 111: 1–41.
- Cappelen, H., and E. Lepore (2003a). 'Radical and Moderate Pragmatics', in Z. Szabo (ed.), *Semantics and Pragmatics* (Oxford: Oxford University Press).
- and — (2003b). 'Unarticulated Constituents and Hidden Indexicals', in M. O'Rourke and C. Washington (eds.), *Essays in Honor of John Perry* (Cambridge, Mass.: MIT Press).
- and — (forthcoming). *Insensitive Semantics* (Cambridge, Mass.: Blackwell).
- Carlson, F., and F. Pelletier., eds (1995). *The Generic Book* (Chicago: University of Chicago Press).
- Carruthers, P. (2002). 'The Cognitive Functions of Language', *Behavioral and Brain Sciences*, 25: 261–316.
- Castañeda, H. (1967). 'Comments', in Rescher (1967).
- Chomsky, N. (1957). *Syntactic Structures* (The Hague: Mouton).
- (1965). *Aspects of the Theory of Syntax* (Cambridge, Mass.: MIT Press).
- (1966). *Cartesian Linguistics* (Lanham, Md.: University Press of America).
- (1969). 'Quine's Empirical Assumptions', in D. Davidson and J. Hintikka (eds.), *Words and Objections: Essays on the Work of W. V. Quine* (Dordrecht: Reidel).
- (1970). 'Remarks on Nominalization', in R. Jacobs and R. Rosenbaum (eds.), *Readings in English Transformational Grammar* (Waltham: Ginn).
- (1977). *Essays on Form and Interpretation* (New York: North Holland).
- (1981). *Lectures on Government and Binding* (Dordrecht: Foris).
- (1986). *Knowledge of Language* (New York: Praeger).
- (1993). 'Explaining Language Use', *Philosophical Topics*, 20: 205–31.
- (1995a). 'Language and Nature', *Mind*, 104: 1–61.
- (1995b). *The Minimalist Program* (Cambridge, Mass.: MIT Press).
- (1996). *Powers and Prospects* (Boston: South End Press).
- (2000a). *New Horizons in the Study of Language and Mind* (Cambridge: Cambridge University Press).
- (2000b). 'Minimalist Inquiries', in R. Martin, D. Michaels, and J. Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik* (Cambridge, Mass.: MIT Press).
- Collins, J. (2003). 'Horwich's Schemata Meet *Syntactic Structures*', *Mind*, 112: 399–432.
- Crain, S., and Pietroski, P. (2001). 'Nature, Nurture, and Universal Grammar', *Linguistics and Philosophy*, 24: 139–86.
- (2002). 'Why Language Acquisition is a Snap', *Linguistic Review*, 19: 163–83.
- Davidson, D. (1967a). 'Truth and Meaning', *Synthese*, 17: 304–23.
- (1967b). 'The Logical Form of Action Sentences', in Rescher (1967).

296 COMPOSITIONALITY, MEANING, AND CONTEXT

- Davidson, D. (1984). *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press).
- (1985). 'Adverbs of Action', in B. Vermazen and M. Hintikka (eds.), *Essays on Davidson: Actions and Events* (Oxford: Clarendon Press).
- (1989). 'What is Present to the Mind?', in J. Brandl and W. Gombocz (eds.), *The Mind of Donald Davidson* (Amsterdam: Rodopi).
- (2001). *Subjective, Intersubjective, Objective* (New York: Clarendon Press).
- Dummett, M. (1975). 'What is a Theory of Meaning?', in S. Guttenplan (ed.), *Mind and Language* (Oxford: Oxford University Press).
- (1976). 'What is a Theory of Meaning (II)?', in Evans and McDowell (1976).
- (1986). 'A Nice Derangement of Epitaphs: Some Comments on Davidson and Hacking', in Lepore (1986).
- Evans, G. (1982). *Varieties of Reference* (Oxford: Oxford University Press).
- and McDowell, J., eds. (1986). *Truth and Meaning* (Oxford: Oxford University Press).
- Fodor, J. (1987). *Psychosemantics* (Cambridge, Mass.: MIT).
- (1998). *Concepts: Where Cognitive Science went Wrong* (New York: Oxford University Press).
- and Lepore, E. (1992). *Holism: A Shopper's Guide* (Oxford: Blackwell).
- and — (1996). 'What Cannot be Evaluated Cannot be Evaluated, and it Cannot be Supervalued Either', *Journal of Philosophy*, 93: 516–35.
- and — (2002). *The Compositionality Papers* (Oxford: Oxford University Press).
- Frege, G. (1879). *Begriffsschrift* (Halle: Louis Nebert). English translation in J. van Heijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931* (Cambridge, Mass.: Harvard University Press, 1967).
- (1892). 'Function and Concept', in P. Geach and M. Black (tr.), *Translations from the Philosophical Writings of Gottlob Frege* (Oxford: Blackwell, 1980).
- George, A., ed. (1989). *Reflections on Chomsky* (Oxford: Blackwell).
- Graff, D. (2000). 'Shifting Sands', *Philosophical Topics*, 28: 45–81.
- Grandy, R. (1973). 'Reference, Meaning, and Belief', *Journal of Philosophy*, 70: 439–52.
- Hauser, M., N. Chomsky, and W. Fitch (2002). 'The Faculty of Language', *Science*, 298: 1569–79.
- Herburger, E. (2001). *What Counts* (Cambridge, Mass.: MIT Press).
- Hermer, L., and E. Spelke (1994). 'A Geometric Process for Spatial Reorientation in Young Children', *Nature*, 370: 57–9.
- and — (1996). 'Modularity and Development: The Case of Spatial Reorientation', *Cognition*, 61: 195–232.
- Hermer-Vazquez, L., E. Spelke, and A. Katsnelson (1999). 'Sources of Flexibility in Human Cognition', *Cognitive Psychology*, 39: 3–36.

- Higginbotham, J. (1983a). 'Logical Form, Binding and Nominals', *Linguistic Inquiry*, 14: 395–420.
- (1983b). 'The Logical Form of Perceptual Reports', *Journal of Philosophy*, 80: 100–27.
- (1985). 'On Semantics', *Linguistic Inquiry*, 16: 547–93.
- (1986). 'Linguistic Theory and Davidson's Program', in Lepore (1986).
- (1989a). 'Knowledge of Reference', in A. George (ed.), *Reflections on Chomsky* (Oxford: Blackwell).
- (1989b). 'Elucidations of Meaning', *Linguistics and Philosophy*, 12: 465–517.
- (1990). 'Contexts, Models, and Meanings', in R. Kempson (ed.), *Mental Representations: The Interface between Language and Reality* (Cambridge: Cambridge University Press).
- (1993). 'Grammatical Form and Logical Form', *Philosophical Perspectives*, 7: 173–96.
- (1994). 'Priorities in the Philosophy of Thought', *Proc. Aristotelian Society* (supp. vol.) 20: 85–106.
- (1998). 'On Higher-Order Logic and Natural Language', *Proc. of the British Academy*, 95: 1–27.
- and May, R. (1981). 'Questions, Quantifiers, and Crossing', *Linguistic Review*, 1: 47–79.
- Hinzen, W. (2002). 'Meaning without Belief', in W. Hinzen and H. Rott (eds.), *Belief in Meaning: Essays at the Interface* (Frankfurt: Hansel-Hohenhausen).
- Hornstein, N. (1984). *Logic as Grammar* (Cambridge, Mass.: MIT Press).
- and D. Lightfoot (1981). *Explanation in Linguistics* (London: Longman).
- Horwich, P. (1997). 'The Composition of Meanings', *Philosophical Review*, 106: 503–32.
- (1998). *Meaning* (Oxford: Oxford University Press).
- Jackendoff, R. (1990). *Semantic Structures* (Cambridge, Mass.: MIT Press).
- (1997). *The Architecture of the Language Faculty* (Cambridge, Mass.: MIT Press).
- (2002). *Foundations of Language* (Oxford: Oxford University Press).
- Kaplan, D. (1989). 'Demonstratives', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (New York: Oxford University Press).
- Katz, J., and Fodor, J. (1963). 'The Structure of a Semantic Theory', *Language*, 39: 170–210.
- Kratzer, A. (1996). 'Severing the External Argument from its Verb', in J. Rooryck and L. Zaring (eds.), *Phrase Structure and the Lexicon* (Dordrecht: Kluwer Academic Publishers).
- Kripke, S. (1971). 'Identity and Necessity', in M. Kunitz (ed.), *Identity and Individuation* (New York: NYU Press).
- (1976). 'Is there a Problem about Substitutional Quantification?', in Evans and McDowell (1976).

298 COMPOSITIONALITY, MEANING, AND CONTEXT

- Kripke, S. (1980). *Naming and Necessity* (Cambridge, Mass.: Harvard University Press).
- (n.d.). 'Reference and Existence: The John Locke Lectures for 1973', unpublished.
- Lakoff, G. (1970). *Irregularity in Syntax* (New York: Holt, Rinehart, & Winston).
- (1987). *Women, Fire, and Dangerous Things* (Chicago: University of Chicago Press).
- Larson, R., P. Ludlow, and M. den Dikken (forthcoming). 'Intentional Transitives'.
- Larson, R. and Segal, (1995), *Knowledge of Meaning* (Cambridge, Mass.: MIT Press).
- Laurence, S., and Margolis, E. (2001). 'The Poverty of Stimulus Argument', *British Journal for the Philosophy of Science*, 52: 217–76.
- Lepore, E. (1983). 'What Model-Theoretic Semantics Cannot Do', *Synthese*, 54: 167–87.
- ed. (1986). *Truth and Interpretation* (Oxford: Blackwell).
- Levin, B., and Rappaport, M. (1995). *Unaccusativity: at the Syntax–Semantics Interface* (Cambridge, Mass.: MIT Press).
- Lewis, D. (1972). 'General Semantics', in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (Dordrecht: Reidel).
- (1986). *On the Plurality of Worlds* (Oxford: Blackwell).
- Longobardi, G. (1994). 'Reference and Proper Names', *Linguistic Inquiry*, 25: 609–65.
- McGee, V., and McLaughlin, B. (1994). 'Distinctions without a Difference', *Southern Journal of Philosophy*, 33: 203–53.
- McGilvray, J. (1998). 'Meanings are Syntactically Individuated and Found in the Head', *Mind and Language*, 13: 225–80.
- (1999). *Chomsky: Language, Mind and Politics* (Cambridge: Polity Press).
- McGinn, C. (1993). *Problems in Philosophy* (Oxford: Blackwell).
- Matthews, R. (2003). 'Does Linguistic Competence Require Knowledge of Language?', in Barber (2003).
- May, R. (1985). *Logical Form: Its Structure and Derivation* (Cambridge, Mass.: MIT Press).
- Meinong, A. (1904). 'Über Gegenstandstheorie', in A. Meinong (ed.), *Untersuchungen zur Gegenstandstheorie und Psychologie* (Leipzig: Barth).
- Montague, R. (1970). 'English as a Formal Language', reprinted in Montague (1974).
- (1974). *Formal Philosophy* (New Haven: Yale University Press).
- Moravcsik, J. (1975). *Understanding Language* (The Hague: Mouton).
- (1998). *Meaning, Creativity, and the Partial Inscrutability of the Human Mind* (Stanford, Calif.: CSLI).
- Neale, S. (1990). *Descriptions* (Cambridge, Mass.: MIT Press).
- (1993). 'Grammatical Form, Logical Form, and Incomplete Symbols', in A. Irvine and G. Wedeking (eds.), *Russell and Analytic Philosophy* (Toronto: University of Toronto).
- Parsons, T. (1990). *Events in the Semantics of English* (Cambridge, Mass.: MIT Press).
- Partee, B., and Rooth, M. (1983). 'Generalized Conjunction and Type Ambiguity', in Bäuerle *et al.* (1993).

- Pietroski, P. (1998). 'Actions, Adjuncts, and Agency', *Mind*, 107: 73–III.
- (2000a). *Causing Actions* (Oxford: Oxford University Press).
- (2000b). 'The Undeformed Domain of Semantics', *Sats, Nordic Journal of Philosophy*, 1: 161–76.
- (2002). 'Function and Concatenation', in Preyer and Peters (2002).
- (2003a). 'The Character of Natural Language Semantics', in Barber (2003).
- (2003b). 'Small Verbs, Complex Events: Analyticity without Synonymy', in L. Antony and N. Hornstein (eds.), *Chomsky and his Critics* (Cambridge, Blackwell).
- (2003c). 'Quantification and Second-Order Monadicity', *Philosophical Perspectives*, 17: 259–98.
- (2004). *Events and Semantic Architecture* (Oxford: Oxford University Press).
- Pustejovsky, J. (1995). *The Generative Lexicon* (Cambridge, MA: MIT Press).
- Putnam, H. (1975). 'The Meaning of "Meaning"', in K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis: University of Minnesota Press).
- Preyer, G., and G. Peters eds. (2002) *Logical Form and Language* (Oxford: Oxford University Press).
- Quine, W. V. O. (1951). 'Two Dogmas of Empiricism', *Philosophical Review*, 60: 20–43.
- (1960). *Word and Object* (Cambridge, Mass.: MIT Press).
- Rescher, N., ed. (1967). *The Logic of Decision and Action* (Pittsburgh: University of Pittsburgh Press).
- Ryle, G. (1949). *The Concept of Mind* (Chicago: University of Chicago Press).
- Sainsbury, M. (1990). 'Concepts without Boundaries', in R. Keefe and P. Smith (eds.), *Vagueness: A Reader* (Cambridge, Mass.: MIT Press).
- Schein, B. (1993). *Events and Plurals* (Cambridge, Mass.: MIT Press).
- (2002). 'Events and the Semantic Content of Thematic Relations', in Preyer and Peters (2002).
- (forthcoming). *Conjunction Reduction Redux* (Cambridge, Mass.: MIT Press).
- Schiffer, S. (1992). 'Belief Ascription', *Journal of Philosophy*, 89: 499–521.
- (1994a). 'The Language-of-Thought Relation and its Implications', *Philosophical Studies*, 76: 263–85.
- (1994b). 'A Paradox of Meaning', *Nous*, 28: 279–324.
- (2000). 'Critical Study: Horwich on Meaning', *Philosophical Studies*, 50: 527–36.
- Soames, S. (1987). 'Direct Reference, Propositional Attitudes, and Semantic Content', *Philosophical Topics*, 15: 47–87.
- (1995). 'Beyond Singular Propositions', *Canadian Journal of Philosophy*, 25: 515–50.
- Spelke, E. (2002). 'Developing Knowledge of Space: Core Systems and New combinations', in S. Kosslyn and A. Galaburda (eds.), *Languages of the Brain* (Cambridge, Mass.: Harvard University Press).
- (2003). 'What Makes Humans Smart?', in D. Gentner and S. Goldin-Meadow (eds.), *Advances in the Investigation of Language and Thought* (Cambridge, Mass.: MIT Press).

300 COMPOSITIONALITY, MEANING, AND CONTEXT

- Stainton, R. (forthcoming). 'Meaning and Reference Some Chomskyan Themes', in E. Lepore and B. Smith (eds.), *Handbook of Philosophy of Language* (Oxford: Oxford University Press).
- Stanley, J. (2000). 'Context and Logical Form', *Linguistics and Philosophy*, 23: 391–424.
- (2002). 'Making it Articulated', *Mind and Language*, 17: 149–68.
- and Szabo, Z. (2000). 'Quantifier Domain Restriction', *Mind and Language*, 15: 219–62.
- Strawson, P. (1950). 'On Referring', *Mind*, 59: 320–44.
- Tarski, A. (1933). 'The Concept of Truth in Formalized Languages', reprinted in Tarski (1983).
- (1983). *Logic, Semantics, Metamathematics*, tr. J. H. Woodger, 2nd edn., ed. J. Corcoran (Indianapolis: Hackett).
- Taylor, B. (1985). *Modes of Occurrence* (Oxford: Blackwell).
- Travis, C. (1985). 'On what is Strictly Speaking True', *Canadian Journal of Philosophy*, 15: 187–229.
- (1996). 'Meaning's Role in Truth', *Mind*, 105: 451–66.
- Uriagereka, J. (2002). *Derivations: Exploring the Dynamics of Syntax* (London: Routledge).
- and M. Piatelli-Palmarini (forthcoming). 'The Immune Syntax', in L Jenkins (ed.), *Variation and Universals in Bilingualistics* (Oxford: Elsevier).
- Vlach, F. (1983). 'On Situation Semantics for Perception', *Synthese*, 54: 129–52.
- Williamson, T. (1994). *Vagueness* (London: Routledge).
- Wittgenstein, G. (1921). *Tractatus Logico-Philosophicus*, tr. D. Pears and B. McGuinness (London: Routledge & Kegan Paul).
- (1953). *Philosophical Investigations* (New York: Macmillan).