

[forthcoming in *The Linguistic Review*: special issue on Jackendoff's *Foundations of Language* and the role of linguistics in cognitive science]

[penultimate version]

The Nature of Semantics: On Jackendoff's Arguments

Steven Gross

Department of Philosophy

University of Pennsylvania

gross2@phil.upenn.edu

The Nature of Semantics: On Jackendoff's Arguments

Abstract

Jackendoff defends a mentalist approach to semantics that investigates conceptual structures in the mind/brain and their interfaces with other structures, including specifically linguistic structures responsible for syntactic and phonological competence. He contrasts this approach with one that seeks to characterize the intentional relations between expressions and objects in the world. The latter, he argues, cannot be reconciled with mentalism. He objects in particular that intentionality cannot be naturalized and that the relevant notion of object is suspect. I critically discuss these objections, arguing in part that Jackendoff's position rests on questionable philosophical assumptions.

The Nature of Semantics: On Jackendoff's Arguments

1. Introduction

We all want meaning in our lives. But what form should our search for meaning take? There is no consensus answer to this question, nor even consensus as to what it asks. Just so in linguistics and cognitive science. Meaning, Ray Jackendoff notes in his *Foundations of Language*, is the “‘holy grail’ not only of linguistics, but also of philosophy, psychology, and neuroscience—not to mention more distant domains such as cultural and literary theory.”¹ But how to pursue this grail? What form should the study of meaning—semantics—take?

Part III of Jackendoff's book continues his development and defense of a *mentalist* approach to semantics. Mentalists explain cognitive capacities in terms of cognitive structures in the mind/brain. These structures, though neurally instantiated, are identified by their cognitive *functions*. Mentalist explanations thus abstract away from details of neural implementation—at least so far as successful explanation requires and allows.² A mentalist approach to *semantics* focuses in particular on *conceptual* structures in the mind/brain and their interfaces with other structures, including specifically

¹ Jackendoff (2002: 267). Henceforth, page numbers in parentheses without further bibliographical information refer to this text. Since the bulk of my remarks are critical, let me emphasize that Jackendoff's book is a major contribution. It's a *tour de force* of synthesis, filled with fascinating hypotheses both large and small. My objections do not affect the parts of the book I do not discuss—in particular, its contributions to lexical and phrasal semantics.

² Mentalism is a variety of what philosophers call 'functionalism'—not to be confused with what linguists label 'functionalism': the attempt to “derive grammatical categories from the exigencies of communication.” (21, fn. 1) To characterize something functionally is to characterize it in terms of its typical causes and effects: a carburetor's functional role is to receive fuel, mix it with the right amount of air, and pass the mixture on. Jackendoff (19-23) reviews mentalism.

Proponents of functional explanations of *cognitive* capacities often talk of mental *representations*—the functionally characterized mental particulars on which the relevant mental processes operate in such explanations. Jackendoff, however, eschews such talk as misleadingly suggesting that mentalists must commit themselves concerning “the thorny philosophical problem of *intentionality*: the apparent ‘aboutness’ of thoughts and other mental entities in relation to the outside world.” (20) See below.

linguistic structures responsible for syntactic and phonological competence. Jackendoff calls this “conceptualist semantics” (271).

Jackendoff contrasts his conceptualist semantics (henceforth, CS) with a “common sense” (294) or “realist” (300) approach according to which the point of semantics is to characterize the intentional relations between expressions and the world, perhaps as mediated by concepts. Semantics, on this alternative approach, would characterize what things expressions *refer* to, under what conditions sentences are *true*, etc.³ Anticipating his main objections to it, I label this ‘intentional worldly semantics’ (IWS).

Now, it’s not obvious that CS and IWS are incompatible. Perhaps they address themselves to different aspects of a larger project. On one common schema, IWS characterizes the semantic facts that a competent speaker of a language cognizes in virtue of the relevant cognitive structures CS characterizes. The semantic facts of my language would include, for example, that ‘water’ refers to water and that (complicating details aside) a sentence of the form [S [NP N] [VP [V]]] is true if and only if what N refers to satisfies V. My semantic *competence* would consist in part in my cognizing these facts—that is, in my having cognitive structures that represent those facts and play a causal role in my linguistic behavior.⁴

But Jackendoff claims that actually IWS “is not so easy to reconcile” with mentalism. (294) Indeed, his arguments, if successful, support a stronger claim—that

³ I bracket (for now) relativization to a conversational context, required to accommodate demonstratives, indexicals, and other phenomena.

⁴ See, for example, Larson and Segal (1995). Note that the intentionality both of linguistic expressions and of concepts is invoked here. The linguistic expression ‘water’ *refers* to water. But there is also a conceptual structure that *represents* this very fact: that ‘water’ refers to water.

There are various ways one might approach this question of reconciling “psychological” and (applied) “mathematical” conceptions of semantics. For some classic discussion, see, for example, Partee (1979), from whom I borrow the quoted labels, and Soames (1985).

IWS is problematic in itself. For they lead him to deny that “conceptual structures [or indeed linguistic expressions or any other candidates] are symbols or representations of anything in the world, that they *mean* anything.” (306, emphasis in original) It can thus seem that Jackendoff denies the very holy grail he seeks. He claims, however, only to deny a mistaken conception of meaning.

In what follows, I critically examine in turn Jackendoff’s two most significant arguments against IWS and thus against the possibility of reconciling it with mentalism: first, that one cannot make “naturalistic sense” (300) of intentionality (roughly, one cannot fit the notion of “aboutness” into a scientific world-view); and, second, that many of the things that expressions and concepts are supposed to be about—and perhaps the IWS *notion* of object itself—are “suspect.”⁵ Whatever the ultimate merits of Jackendoff’s *conclusions*, his arguments are unconvincing. Their point and basis are not always clear; and, when they are, they commit the defender of CS to controversial, but undefended, theses. Moreover, implicit in these arguments are *philosophical* assumptions in tension with a *properly* “naturalistic” attitude. Uncovering these assumptions indicates a better, though less conclusive, strategy for defending conclusions like Jackendoff’s.

2. Intentionality

⁵ Jackendoff (Chap. 10) frames his challenge a little differently. He claims that there are problems with the IWS conception of *language* and with its notion of object. Of the problems he raises for the IWS conception of language, however, only the problem of “naturalizing” intentionality concerns something essential to IWS. For example, Jackendoff objects to versions of IWS that treat language as something “out there in the world” instead of in “the f[unctional]-mind of language users.” (296-7) But some IWS theorists *agree* that the proper objects of empirical study are internal languages—I-languages, in Chomsky’s (1986) sense. Cf. Larson and Segal (1995). Also, Jackendoff (297-300) raises a problem specifically for views that consider languages abstract objects: since abstract objects are by definition non-material, non-spatio-temporal objects that cannot stand in causal relations, relations to abstract objects cannot be “naturalized”; thus nothing about the mind/brain of a speaker could make it the case that a particular language was hers. I do not believe the conclusion in fact follows, but in any event proponents of IWS need not maintain that languages are abstract objects. (Incidentally, Jackendoff underplays the fact that nominalism—the rejection of abstract objects—faces well-known obstacles of its own. See, for example, Burgess and Rosen (1997).)

Jackendoff's brief against intentionality takes a form familiar to philosophers. "There is no magic," he writes, "... we seek a thoroughly naturalistic explanation that ultimately can be embedded in our understanding of the physical world." (268) But "one cannot make naturalistic sense of intentionality." (300) There is thus no such thing as intentionality: "there is ... no physically realizable causal connection between concepts and objects"—at least no causal connection that could determine intentional relations. (300)⁶ The IWS conception of semantics, which gives pride of place to intentionality, is therefore ill-begotten.

To assess this argument we need to clarify what making naturalistic sense of something requires—and why Jackendoff holds that intentionality does not make the grade. It will be useful first, however, to underscore how controversial Jackendoff's conclusion is.

2.1. *For intentionality*

While not a view without its supporters, most famously Quine,⁷ the denial of intentionality runs counter to the views of a large number of semanticists working within

⁶ I add the qualification because otherwise the claim seems obviously false: surely, the presence of a cow can activate my cow-concept. It's possible—though I think unlikely—that Jackendoff would respond that cows do not exist because they would be problematic objects. (For instance, their temporal and spatial boundaries are vague. See below.) But in any event surely concepts interact with *some* objects, even if only with other neurally instantiated cognitive structures.

⁷ See Quine (1960). Cf. Churchland (1981); and Stich (1983). Churchland's claim, however, is not that there's no intentionality, but that what intentionality there is, is unlike that ascribed in everyday discourse. Cf. Churchland (1989). The denial of intentionality is often labeled 'eliminativism.' Jackendoff perhaps intends to distance himself from some forms of eliminativism when he stresses (295) that the commonsense view embodied in IWS "cannot be eliminated from everyday discourse." Cf. Jackendoff (1991: 424), citing Stich and Churchland's views as too "brittle". But the crucial eliminativist claim concerns what there is, not what people will *speak* as if there is. Cf. Churchland (1981: 85-6). Other remarks by Jackendoff likewise indicate an interest in qualification: for instance, that his "goal is not to show that the common-sense view of reference and reference and truth is false and should be rejected out of hand." (294) It's hard to square these remarks with his actual line of argument and the quotes above. (Perhaps he thinks the common-sense

a variety of frameworks—as Jackendoff emphasizes. Independent of this sociological fact, there are substantive reasons for demurring.

First, the existence of intentionality can seem evident from one’s own case. As I type, I am currently thinking about Jackendoff—for example, how he might best respond to my remarks. I am also thinking about the flashing light on my mouse—how it indicates that the batteries need recharging. I am more certain about the fact that I’m thinking such thoughts *about* such things than I am about most other topics. Certainly, I am more certain about them than I am about how best to pursue semantics. What better evidence for intentionality could there be? Call this the *argument from self-evidence*.⁸

Second, some maintain that to deny intentionality is incoherent. In its crudest form, the objection is that those who deny intentionality inconsistently believe that there are no beliefs. Deniers of intentionality may reply that people merely find it useful—not to mention unavoidable—to take up an “intentional stance” towards others, to act *as if* their utterances exhibit intentionality. But it’s claimed that this reply forgets that *taking up the intentional stance*—and claiming and denying, for that matter—are themselves examples of intentionality. Call this the *argument from incoherence*.⁹

Third, some maintain that there is significant empirical research that essentially employs intentional vocabulary—not only in semantics, but also for example in the study of vision and conceptual development. It is claimed that the fact of *successful* science employing intentional vocabulary provides strong support for the naturalistic legitimacy

is neither true nor false? Perhaps he only rejects this characterization of his goal because he rejects the “should be rejected out of hand” part?) I concentrate on those elements of his book that deny intentionality.

⁸ The claim of self-evidence is critically discussed by Wright (2002: 205) and by Rey (1997: 83-8).

⁹ There are replies to these initial formulations. For a sophisticated development of the incoherence worry, see Wright (2002). The phrase “intentional stance” comes from Dennett (1971).

of the phenomena which that vocabulary purports to denote. Call this the *argument from successful science*.¹⁰

Fourth, there are those who take themselves to have offered positive proposals for relating intentional properties to other properties that are not controversial from a scientific point of view. For example, some maintain that strategies exist that enable one in principle to *identify* intentional properties with *physical* properties. Even where these strategies (and others—we will mention an example in a moment) are really barely more than strategy-sketches, it is maintained that sufficient direction has been provided for us at least to envision what successful naturalization of this sort might look like. Call this the *argument from physicalism*.¹¹

A full defense of the denial of intentionality would have to adequately respond to these four reasons in its favor. A denial that turns on intentionality's *naturalistic* prospects must, in particular, address *both* the argument from successful science and the argument from physicalism. Jackendoff, however, concentrates on one version of the argument from physicalism. Having taken himself to have established the non-naturalizability of intentionality and thus, on his view, its non-existence, he attempts “to show why the common-sense view is nevertheless so intuitively compelling,” (295) an effort that can be construed as an attempt to explain away the intuitions behind the argument from self-evidence—one would hope without falling prey to the argument from incoherence.

¹⁰ Ironically, philosophers often point to the success of linguistics as an example of the scientific legitimacy of intentionality. Thus Rey (1997, Chap. 4) deploys Chomsky against behaviorist eliminativists (behaviorists who, not only prescind from mentalistic vocabulary, but also deny the existence of such phenomena as intentionality and consciousness). However, Chomsky disavows intentionalist readings of his work. Cf. Rey (2003a); Chomsky (2003b) in reply; and Rey (2003b), a rejoinder. Burge (2003) adverts to vision research in reply to Chomsky (2000, Chap. 7).

¹¹ Loewer (1997) is a survey of such attempts to naturalize intentionality. Stich and Warfield (1994) is a representative collection of papers.

I will not comment on Jackendoff's attempt to explain away our intentional talk—he acknowledges that it stands in some need of development (332)—beyond noting that it in fact displays how difficult it can be to avoid intentional vocabulary in characterizing and explaining cognitive phenomena.¹² I will focus rather on Jackendoff's case for the non-naturalizability and thus non-existence of intentionality. Jackendoff's reply to the argument from physicalism is weak. But more importantly it implicitly endorses a problematic conception of naturalism that underplays arguments from successful science—precisely the arguments on which the status of intentionality should turn, according to a less problematic conception of naturalism. Moreover, when naturalizability is properly construed, the *non*-naturalizability of intentionality does not entail the *denial* of intentionality. A successful argument for the non-naturalizability of intentionality would then not require a response to the arguments from self-evidence and from incoherence. Or so I shall suggest.

2.2 Jackendoff's naturalism and his criticism of Fodor

What then does Jackendoff mean by “making naturalistic sense” of a phenomenon? Above, I glossed such talk in terms finding a place in a scientific world-view, a scientific conception of or approach to nature. But this leaves room for quite distinct elaborations. Jackendoff's remarks offer very little to go on, but we have already seen that he twice alludes to a privileged status for *physics*. This suggests that the naturalism he favors is

¹² Jackendoff, for example, aims to replace the “common sense realist theory of reference: phrase P in language L, uttered in context C, refers to entity E in the world (or possible worlds)” with “a conceptualist theory of reference: a speaker S of language L judges phrase P, uttered in context C to refer to entity E in [the world as conceptualized by S]”. (304—Jackendoff's brackets) But taken at face value this says that speakers bear certain attitudes—judging—to certain intentional contents. Jackendoff's further explication of “judges” (304, fn. 7) also employs seemingly intentional vocabulary.

some version of *physicalism*—the doctrine that everything there is, is constituted of physical things (the entities posited by physics), and all their properties are, or are properly related to, physical properties (the properties posited by physics).

Physicalism itself, however, comes in various strengths, depending on what constitutes a proper relation to physical properties. *Reductive* physicalism requires that a naturalistically acceptable property be *identical* to a (perhaps highly complex) physical property. A paradigmatic naturalization might then be the reduction of temperature to molecular mean kinetic energy. *Non-reductive* physicalism does not require identity. For instance, on some versions it suffices that the putative property so depend—or, “supervene”—on physical properties that any change in the non-physical property must involve a change in underlying physical properties, but not necessarily vice versa. (Compare: a change in my screen’s global picture-properties—for instance, being symmetric—requires a change in pixel-properties, but not necessarily vice versa. See also fn. 14 below.)¹³

One can speculate that Jackendoff would favor some version of *non-reductive* physicalism. First, it comports better with his own defense of mentalism (23) against those who maintain that a neuroscience purged of cognitive vocabulary can suffice to explain our cognitive capacities. Second, it’s suggested by his *argument* against the naturalizability of intentionality. For Jackendoff’s argument consists in a criticism of Jerry Fodor’s attempts to “develop a ‘naturalized’ semantics, that is ‘to say in *nonsemantic* and *nonintentional* . . . terms what makes something a symbol.’” (279, quoting Fodor (1990b), incorrectly cited as (1991)) But Fodor is famously a non-reductive physicalist: the non-semantic, non-intentional terms need not be themselves in

¹³ Stoljar (2001) surveys varieties of physicalism and standard arguments for and against.

principle *reducible* to the language of physics.¹⁴ Jackendoff denies that Fodor’s attempt to naturalize intentionality is successful, but not by challenging his conception of the goal.

What, then, is Jackendoff’s criticism? Fodor (1990a, 1990b) advances a version of the “causal/covariational,” or “informational,” approach to naturalizing intentionality. Such views start from the idea that one’s concepts tend to be caused by what they are about. The presence of a nearby cow, for example, will cause the tokening of one’s cow-concept. We might thus think of symbols as roughly causally covarying with their referents as a thermometer’s mercury level covaries with—and thus represents—temperature. However, one can also think about cows in the absence of cows, for example in the presence of a well-disguised moose or a glass of milk. So, the content of one’s concepts cannot be equated with what causes it to be tokened. One needs somehow to weed out the causes irrelevant to the fixing of content. Fodor’s idea, roughly, is that the content of a concept—what it’s about—is the cause upon which its *other* causes depend. That cows tend to cause tokenings of one’s cow-concept does not depend on the fact that a well-disguised moose would as well. But a well-disguised moose would

¹⁴ Fodor (1974, 1998b) argues, following Putnam (1967), that many explanatorily relevant properties may be physically realizable in a *multiplicity* of ways and thus not *reducible* to physical properties. Indeed, he holds that this is the case with many functionally characterized *mental* properties. Believing that $2+2=4$ or representing the presence of an edge in the visual field may be realized, for example, in *one* kind of physical stuff in humans (ensembles of neurons)—but we cannot rule out in advance their being realized in *different* kinds of physical stuff in other beings or machines. Further, Fodor’s goal—see (1990a: 96)—is only to state, in non-semantic, non-intentional terms, conditions *sufficient* for having intentional properties. So, even if *these* conditions *are* reducible to physical conditions, it does not follow that the intentional properties are. For that, one would need reducible *necessary* and sufficient conditions. (*Cognoscenti* will recognize that here and elsewhere I am running roughshod over a variety of important distinctions—for example, between analyzing a *predicate* in terms of others vs. reducing a *property* to others. They do not matter for the points I wish to make.)

cause tokenings of one's cow-concept only because cows would as well. Jackendoff (279) objects that it's left unexplained *why* this dependency obtains.¹⁵

3.3. *Assessing the argument*

We are now in a position to assess Jackendoff's basis for denying intentionality. I proceed roughly in order of increasing importance. First, Jackendoff's criticism of Fodor—his sole reason for denying the naturalizability of intentionality—does not cut very deeply. Jackendoff does not question the plausibility of Fodor's dependency claim. He only asks for an explanation of its obtaining. But it's unclear why Fodor must explain *its* basis in order to have succeeded in his goal of stating in non-intentional terms sufficient conditions for having intentional properties. Perhaps Jackendoff suspects that one cannot cash out the dependency claim itself in a physicalistically acceptable way or in a way that does not advert to intentionality. But, even if this speculation is correct, Jackendoff would need to supply some supporting argument. In sum, the criticism is just insufficiently developed to have much force as it stands.¹⁶

Second, with all due respect to Fodor, that he has not shown us how to do something would not show that it cannot be done. Indeed, there are other physicalist strategies on the market for naturalizing intentionality—for instance, teleological strategies that analyze intentionality in terms of a cognitive structure's evolved and/or

¹⁵ Jackendoff's confidence is bolstered by various expressions of uncertainty he finds in Fodor. The second Fodorian aside he quotes (279, fn. 6) actually concerns a different issue. But in any event Fodor does not himself hold that the naturalizability of intentionality rests on the success of his own suggestion. Jackendoff also parenthetically criticizes Fodor for not saying *how* something can cause a representation of it to be tokened: by "acting on the speaker's perceptual system? Fodor doesn't say." (279) This seems unfair, since the answer is obviously 'yes'.

¹⁶ Fodor (1990a) anticipates and discusses Jackendoff's objection and many others. My point, however, is not to defend the position, but only to note that Jackendoff does not supply much by way of a developed objection to it.

learned function, and inferential role strategies that analyze intentionality in terms of causal/computational role. Their details need not detain us here.¹⁷ But to show that intentionality cannot be naturalized, one must address versions of the argument from physicalism that advert to these alternative attempts.

Third, even this would of course only shift the burden of proof to the other side; for it would not follow, if problems were identified with extant strategies, that no better strategy was forthcoming. Indeed, fourth, even if it were conceded that we will not and *cannot* construct a compelling argument from physicalism, it would not follow that intentional properties are not physicalistically acceptable (and thus not in this sense naturalizable and thus not instantiated). For our inability to *demonstrate* that intentional properties bear the appropriate relations to physical properties might merely reflect our own cognitive limitations. The world might not be limited to what evolutionary pressures

¹⁷ Cf., respectively, Millikan (1989); and Block (1986), both reprinted in Stich and Warfield (1994). Some of Jackendoff's remarks suggest some sympathy for inferential role strategies:

Fodor (1998a) objects to [Jackendoff's proposals concerning lexical] decomposition, on the grounds that we have gotten no closer to explicating meaning if the meanings of the primitives have not themselves been explained My working hypothesis is that the meaning of the primitives is exhausted by (a) their capacity for combination with other primitives and (b) their use in triggering inference rules and interface rules, individually and in combination. But I acknowledge that this remains to be demonstrated. (369)

It might seem strange that Jackendoff, while denying intentionality, accedes to an objection concerning the primitives' intentional content. The explanation is that Jackendoff does not use the term 'meaning' as Fodor does. For Jackendoff, meanings are "pure non-intentional structure" (279) that do what "meaning is supposed to do, such as support inference and judgment." (306) (Again, talk of judgment must itself ultimately be purged or supplied a non-intentional explication.) Jackendoff is gesturing at what he thinks the functional role of primitive concepts are. But he does not think that to specify this functional role is to specify what bestows intentional content (there being none). Thus Fodor and Jackendoff are here two ships passing in the night.

A terminological remark. Strategies for naturalizing intentionality are often labeled 'semantics': informational semantics, teleo-semantics, inferential-role semantics, as the case may be. But these are not themselves instances of IWS. We must distinguish *foundational* and *descriptive* semantic projects. The former attempt to provide a foundation for the latter (IWS) by showing how the properties IWS describes are related to non-intentional properties. It is easier to recognize the distinction if one keeps in mind that naturalizing projects focus on the non-intentional properties in virtue of which *primitive* (atomic) representations have their intentional properties. It is left to a compositional IWS to describe how *complex* representations inherit their intentional properties from their constituents and their mode of combination. Cf. Kaplan (1989) and Stalnaker (1997).

happened to have enabled us to make full sense of.¹⁸ To argue that it's not just *our* limitations that make naturalizing intentionality difficult, Jackendoff would need to identify some aspect of *intentionality* that precludes its bearing the appropriate relations to the physical.¹⁹ This in turn would require a more fully developed conception of non-reductive physicalism—in particular, some more specific characterization of what would count as a property's bearing a naturalistically *appropriate* relation to physical properties.

This brings us to my fifth and final worry: should one accept *any* variety of physicalism in the first place? Chomsky urges that we should not: physicalism is either empty or false—empty, if some unknown future physics is meant, and false, if current physics is meant (since we have no reason to think our conception of the physical will not continue to evolve as new phenomena are explored).²⁰ If this is right, it need not follow that *naturalism* is empty or false. There might be some alternative construal. Chomsky's remarks suggest that we understand naturalism as a *methodological* doctrine: one ought to aim for well-supported, intelligible explanations as judged against the evolving standards of successful science.

Methodological naturalism contrasts with physicalism in three related ways, all of which stem from physicalism's being a *metaphysical* doctrine—one that makes a claim about *how things are*. First, if a phenomenon fails to satisfy *physicalist* constraints, one

¹⁸ Chomsky (1975, Chap. 4) famously refers to mysteries currently beyond our reach, perhaps owing to aspects of our cognitive faculties (performance limitations aside). Maybe aspects of the relation of intentionality to physical phenomena constitute such a mystery. Cf. Chomsky (2000: 45). Jackendoff (423) at one point speaks of “the mystery of intentionality.” I do not know whether he is alluding to Chomsky's usage. My suggestion, however, is not that *intentionality* is a mystery—only that its relation to the physical might in part be.

¹⁹ One common strategy is to argue that intentional properties are in some sense intrinsically *normative* (concerning how things ought to be or how people ought to act) while physical properties are intrinsically *non-normative*. See, for example, Kripke (1982).

²⁰ See, e.g., Chomsky (2000: 143-5 and 117) and (2003a). The latter is a reply to Poland (2003), an overview of Chomsky's challenge to physicalism. Others have also pressed these and related points—for example, Crane and Mellor (1990).

must deny that this phenomenon exists at all. But it does not follow from *methodological* naturalism that, if a putative phenomenon cannot be brought within the ambit of science, then there is no such thing. It only follows that there is a limit to how far we can succeed in satisfying our explanatory ambitions.²¹ This is why a *methodological* naturalist who denies that one can make naturalistic sense of intentionality can consistently accept the arguments from self-evidence and incoherence in favor of the existence of intentionality.

Second, physicalism asserts that the world is unified in a certain way—in whatever way properties must be appropriately related to physical properties. Methodological naturalism makes no such assumption. It *does*, however, have us *aspire* to unification; for unification affords deeper and more broadly applicable explanations. A methodological naturalist might thus *aspire* to relate intentional properties to physical properties without holding hypotheses that advert to intentionality hostage to the fortunes of that aspiration. Note that this applies whether those hypotheses advert to intentionality as the *object* of explanation or in explaining *other* phenomena.²²

Third, we have already noted that nothing in physicalism—at least pending further elaboration—guarantees that physicalistically acceptable properties are such that we can bring them under scientific control. What there is, is not identical with what we

²¹ Cf. Chomsky (2000: 77). Note that a phenomenon's falling outside the scope of our scientific efforts does not preclude our being able to know a fair bit about it. Our cognitive endowment may enable us to know about some things without enabling us to construct successful explanatory theories about them. Cf. Burge (2003) on the role of reflective judgment in semantics.

²² On explanation and unification, see Friedman (1974) and Kitcher (1981). Adapting Kantian jargon, we might say that the physicalist views unification as a *constitutive* principle, while the methodological naturalist views unification—like Kant—as a *regulative* principle.

Jackendoff, rightly defending the need for CS even if it turns out that he is wrong and we must also invoke intentionality (see below), rightly remarks that “there is no reason to be paralyzed by the absence of a solution for intentionality, as Fodor seems to be.” (280) He means that we can independently pursue the non-intentional aspects of semantics. My point is a bit different: the absence of a physicalist “solution for intentionality” does not preclude pursuing scientific projects that *do* invoke intentionality. (I would add that Fodor also agrees with Jackendoff: though he *himself* has chosen to concentrate on, among other things, a very hard problem others sometime ignore, he does not think work such as Jackendoff's must be put on hold in the meanwhile.)

can understand and explain. But on the methodological construal, a naturalistically unacceptable property is precisely one that should not appear in a scientific theory as judged by current standards. And this is why, from the perspective of methodological naturalism, what *really* matters here are arguments from successful science. For, though methodological naturalism does not in itself provide the basis for *denying* intentionality, it certainly allows—indeed, insists—that one reject specific theories of, or adverting to, intentionality that are explanatorily vacuous, unintelligible, insufficiently fruitful, etc.²³ This is not the form that Jackendoff’s criticism takes. But I am suggesting that this where the action should be.²⁴

Jackendoff thus fails to make good on his first objection. So far, then, IWS theorists have been given no reason to think that their project is intrinsically problematic or in tension with mentalism. Nor does this threaten Jackendoff’s CS: its legitimacy does not depend on deligitimizing IWS.²⁵ Whether IWS constitutes interesting science is another—and altogether more important—question. But, whether or not it does, if there *is* intentionality, then the holy grail of meaning cannot be reached by CS alone.²⁶

²³ Chomsky’s hostility to forms of IWS, or at least certain ways of construing IWS, is best understood in this light. Cf. Chomsky (2003c).

²⁴ Jackendoff does remark at one point: “as we will see, [Fodor’s] conclusions are at such odds with all detailed empirical work on meaning as to discredit the enterprise in the eyes of practical semanticists.” (269) It’s possible that this refers forward to Jackendoff’s discussion of Fodor’s attempt at naturalization. But whatever the merits of Fodor’s attempt, Jackendoff nowhere shows that it’s at odds with empirical work in semantics. Perhaps the quote refers rather to Jackendoff’s discussion of other positions Fodor holds, such as his views on lexical decomposition (334-7). This would not be immediately germane to the present point.

²⁵ Jackendoff remarks that, even if Fodor is right and “conceptual structure might indeed need to be intentional in some sense . . . we still have to work out the details of the combinatorial system constituting semantic/conceptual structure/LoT [Language of Thought], as well as its interfaces with language, inference, perception, and action—which is what I take as the task of conceptualist semantics.” (279)

²⁶ I should briefly comment on two other arguments against intentionality one might discern in Jackendoff. First, Jackendoff (19-20) mentions that, if representations must represent something *to* someone, then there is a problem with representations to which subjects have no conscious access (for example, many of the cognitive structures posited by linguists if they were deemed representational). It’s unclear, however, that Jackendoff *endorses* this worry: he merely mentions it as an issue he avoids by eschewing intentional talk. Nor does he bring it up again, even when later challenging the naturalizability of intentionality. In any

3. Objects

Jackendoff's second objection to IWS concerns the things (objects, entities, events, states-of-affair, etc.) that linguistic expressions or conceptual structures are allegedly about. There are two ways of understanding this objection. Judging from the bulk of Jackendoff's remarks, the problem is supposed to be that IWS commits theorists or speakers to maintaining the existence of objects that are in fact problematic in some way: because they clearly do not exist (unicorns (301)), or because their "ontological status is far less clear" than "middle-sized perceivable physical objects" (303), or because they are imperceptible or abstract or dependent upon us for their existence. However, other remarks suggest that Jackendoff's fundamental complaint concerns, not the alleged *objects*, but rather the very *notion* of object IWS theorists deploy, as when he recommends critically "examining the realist's notion of 'objects in the world.'" (300—cf. 303)

event, if this worry *were* pressed, the defender of intentionality could deny that a representation must be a representation *to someone* in a sense that entails accessibility to consciousness: it might suffice that the representation have an appropriate functional role in the subject's cognitive economy. Moreover, the relevant notion of accessibility to consciousness is unclear. Jackendoff suggests a gloss in terms of *actual* access, perhaps with sufficient introspective effort. But even Searle (1990), arguably the most fervent defender of the claim that accessibility to consciousness is a necessary condition on being an intentional mental state, allows that there can be intentional mental states *accessible* to consciousness that are *in fact* beyond the reach of introspection owing to some interfering factor. Indeed, Block (1990) uses this concession to question whether the notion of accessibility to consciousness can be compellingly clarified in a way that excludes the deliverances of linguistic theory: without further clarification, the most abstruse aspects of syntax might in fact qualify as *accessible* to consciousness.

Second, some might discern an argument against intentionality in Jackendoff's reminders that neural and cognitive structures are "trapped in our brains" (306, 423, cf. 305). Whatever the intent of these remarks, however, it's more charitable to read them as stylistic variants of the charge that intentionality—which would relate concepts to things outside the brain—cannot be naturalized. For it obviously does not follow, from the fact that something is in the brain, that it cannot stand in relations to things outside the brain. My being stuck in Washington does not preclude my standing in the relation 'biological brother of' to someone stuck in Baltimore. That this relation supervenes on complex causal relations does not alter the fact. (Some of Jackendoff's remarks suggest that he thinks that intentionality must involve a *direct* causal relation. But I see no reason to preclude an expression or concept's being about something in virtue of its standing in a complex, mediated causal relation to that thing.)

These two ways of developing the objection can seem mutually exclusive: if one challenges the *notion* of object, one cannot *deploy* that notion in assessing whether one's opponent has got its extension right. But this depends on how we understand Jackendoff's informal talk of a *notion*. To a philosopher's ears, it sounds like he is challenging the coherence, intelligibility, or fruitfulness of a *concept*. But perhaps he means only to be challenging a conception of the concept's extension—that is, perhaps he is just again raising the issue of what objects we should in fact affirm. To cover my bases, I will remark on the objection understood both ways: first, as an objection to IWS's problematic *objects*, and then more briefly as an objection to the IWS *notion* (concept) of object.

3.1. *Problematic objects*

Jackendoff proceeds by presenting a variety of well-known problem cases—fictional characters, abstract objects, social entities, etc.²⁷ There is not space to discuss them all here. His examples of “geographical objects” will serve for the comments I wish to make:

There is nothing tangible about Wyoming, no great geographical features that mark it off, no lines drawn across the landscape It's a purely politically constructed entity, its rectilinear boundaries fixed by a stipulative act.

We can touch the Mississippi River, and swim in it. But is the river the water contained in it, the bed of the river, the complex of the two?

²⁷ As Jackendoff notes (303), these cases are much discussed at least in the philosophical literature. (On the same page, however, he criticizes proponents of IWS for deploying an “unexamined” notion of “objects in the world.”) It should not be assumed that these heterogeneous cases raise issues of the same kind.

Exactly where does it end in the Gulf of Mexico, and exactly where does its tributary, the Missouri, end in it? One can draw arbitrary lines on a map, but these are understood as matters of convenience and not some sort of “natural truth about the world.”

The distance between New York and Boston is not tangible. Nor is there an absolute truth about it: how should it be measured? From center to center, from nearest border to nearest border, along some particular highway, on a straight line through the earth’s crust? Much depends on one’s purpose. (301)

Jackendoff’s complaint would seem to be that there is something problematic about these putative objects—and this is a problem *for IWS* because IWS commits theorists or speakers to their existence: for presumably, according to IWS, these are the objects to which linguistic expressions refer. I will argue, however, that, first, Jackendoff moves too quickly in assuming that IWS commits anyone to the existence of *any* objects; second, even if IWS does generate commitments concerning existence (“ontological” commitments), it’s unclear at least in some cases that IWS commits anyone to the *particular* objects Jackendoff discusses; and, third, it’s any case unclear what’s *problematic* about (at least some of) these objects.

3.1.1. No commitment

First, intentional talk—talk of aboutness, reference, and the like—is ambiguous between senses that entail existence and those that do not. In the former sense, neither a speaker

nor her words or concepts can succeed in referring to unicorns, since there are none. But of course there is also a sense in which one *can* think about and refer to unicorns—as I have done in using this very sentence. It's thus not obvious that such candidate IWS claims as that 'unicorn' refers to unicorns must be understood as entailing the existence of unicorns. And if someone were to demand or stipulate an existence-entailing reading, it's not obvious that IWS theorists could not restate such claims accordingly: to wit, 'unicorn' refers to unicorns, if there are any.

From this follows a second point. Typically, one cannot derive an existence claim from a semantic claim without a supplementary *non*-semantic premise. Consider the “eventish” semantic claim that [_s Emma ran quickly] is true if and only if there was something that was a running, of which Emma was the Agent, and it was quick. Does this commit anyone to the existence of events—or at least runnings? Not unless it is also believed that [_s Emma ran quickly] *is* true. But of course it's no part of *IWS* to claim *that*.

In fact, even *with* the requisite supplementary non-semantic premise in place (claim P in a biconditional of the form 'P if and only if Q,' as in the previous paragraph's example), there's room to question whether any ontological commitment follows, either for the theorist or for the speaker. Consider first the theorist. On the view that would reconcile IWS and CS, semantic claims concern *what a particular speaker cognizes*. But if semantic claims are understood as falling within the scope of an attitude ascription, then the IWS *theorist* is not thereby committed to anything beyond the ascription itself. (When I assert that Emma believes in the tooth-fairy, I commit *myself* only to that ascription, not to the existence of the tooth-fairy—similarly for my ascriptions concerning what her words mean.) But then from the fact that the *theorist* believes that P

and also ascribes to the *speaker* the cognition that ‘P if and only if Q,’ you cannot conclude anything about the *theorist’s* attitude towards Q. Now consider the speaker. Recall that the ascriptions largely concern tacit states not under the speaker’s full control—what she “cognizes,” not necessarily what she explicitly believes or knows. She is therefore not *responsible* for their contents in a way that she is for what she *believes*. Thus, it’s a considerable step to conclude, from her *believing* P and her *cognizing* that ‘P if and only if Q,’ that she believes or is otherwise committed to Q, including whatever existence-claims it makes.²⁸

3.1.2. Not *that* commitment

The first reply questioned whether IWS generates any existence claims at all. Many IWS theorists, however, would readily concede—even insist—that it does. But even *they* need not be concerned by Jackendoff’s remarks. Presented with such cases, they might deny that IWS is committed to the existence of the *particular* supposedly problematic object.

Consider Jackendoff’s example of the distance between New York and Boston. This putative object should only concern the IWS theorist if her semantics commits her to its existence. But does it? One’s best guess is that Jackendoff assumes that it does because he assumes that IWS must hold that the phrase ‘the distance between New York and Boston’ refers to it. But that assumption requires support. Indeed, proponents of IWS might question the assumption that the phrase refers to an object at all. On a Russellian view, it’s rather a quantifier phrase, so that sentences containing it make cardinality claims. Still, Jackendoff can object that such sentences as ‘The distance between New

²⁸ I expand on these and related matters in Gross (forthcoming-a). These remarks may seem close to Jackendoff’s conceptualist conception of reference. (Cf. fn. 12 above.) But I view the attitude ascriptions as intentional and abstain from drawing ontological conclusions. More on this below.

York and Boston is about 200 miles' quantify over—and thus commit those who accept the sentence to—the existence of distances, even if the sentence contains no phrase that *refers* to them. And what “in the world” are *distances*? We will examine how one might answer this question in a moment. The point I want to make here, however, is that it's not obvious that the sentence, if true, *does* quantify over distances. As those wary of admitting distances into their ontology will note, it's often proposed that such sentences are really about *numbers*. The truth-condition of the sample sentence would be something like: there's a number n such that $\text{MILES}(\text{New York, Boston}) = n$, and n is equal to a number near 200 (where 'MILES' denotes or expresses a function from pairs of places to numbers). If this view can be defended, then at least *this* sentence raises no *further* issues about distances above and beyond issues that arise about numbers and functions.²⁹

These observations remind us that questions of reference and ontological commitment arise, not for a phrase alone, but for a phrase (1) within a sentential context that's (2) represented in its semantically relevant way. Proponents of IWS need not assume—nor should their critics—that such matters can be easily read off surface form.³⁰ Jackendoff's remark about “absolute truth” and “one's purpose” illustrate, further, that questions of reference and ontological commitment typically arise only for expressions as used in some *conversational* context. An improved account of the sample sentence's truth-conditions might thus add a parameter for, or otherwise accommodate, the

²⁹ Pro-measure expressions, as in Jackendoff's example 'The fish that got away was *this* long' (316), would be understood as contextually introducing a unit of measure. Ironically, Jackendoff's anti-syntactocentrism (107) might make it slightly easier to defend such treatments, since it lessens the demand for independent syntactic evidence for a semantic hypothesis. I do not mean to *assume*, however, that the hypothesis *can* be defended. But by the same token it cannot just be assumed that the sentence quantifies over distances. (After I wrote this paper, Higginbotham (2004) appeared suggesting the very same treatment. Jackendoff (2004) replies.)

³⁰ Cf. the discussion of apparent reference to flaws in Ludlow (2003); and Chomsky (2003c).

contribution of conversational context in determining how the distance between the two cities is to be measured. But this raises no issue about the *objects* thereby referred to.³¹

3.1.3. No problem

The IWS theorist, however, need not rest her case on the hope that any problematic ontological commitment will be analyzed away. She can also ask whether these objects are really problematic in the first place.

Jackendoff raises two problems for the distance between New York and Boston: it's not tangible, and there's no "absolute" truth about it. The latter was just covered. About the former, one might reply, first, that the object is just the spatial extent between those two cities: if one can touch a location in space, then one can touch that. (Or at least a part of it—one might not have the patience to touch it *all*.) But even if one can't, so what? Why must all objects be tangible? It may well be that in its core non-technical use 'object' tends to refer to *perceivable* objects. But proponents of IWS need not stick to this sense of 'object.' Like any theorist, they may find themselves deploying concepts that build but depart from our pre-reflective uses. For example, distinctions of type aside, IWS may deploy the so-called *logical* concept of object, according to which an object is whatever exists.³² Surely, it's at least as reasonable to conclude that there are intangible objects, given the existence of extents of space (supposing they are not tangible), as it is to conclude that there are no extents of space because they would not be tangible.

³¹ It may, however, provide the basis for an empirical challenge to specific ways of attempting to accommodate the role of conversational context in IWS. Cf., for example, Gross (forthcoming-b) and Pietroski (2003). Perhaps one might even challenge the ability of *any* version of IWS to accommodate these facts. Jackendoff holds that it's "important to factor out the respective contributions to understanding made by linguistic expressions and by context." (280) But this is not the basis of his challenge to IWS. Such a challenge would indeed be just the sort a methodological naturalist should take seriously.

³² Cf. Parsons (1982).

Jackendoff notes that Wyoming “is a purely politically constructed entity, its rectilinear boundaries fixed by a stipulative act.” But he does not explain why this is problematic either. A constructed entity is still an entity. It did not exist before the relevant political activities, but it exists now. It can thus stand in relations to expressions. Jackendoff appears to hold that a notion of object that countenances objects that owe their existence to us in some manner is for that reason “suspect.” (303) But why should one think that? Artifacts, such as chairs, owe their existence to us, but it’s unclear that this renders suspect a notion of object that comprises them. It’s true that the existence of agents is a necessary condition for the *continued* existence of Wyoming, but not for the continued existence of a chair. But an argument is needed to show that the difference renders Wyoming (or a notion of object comprising Wyoming) problematic. It’s also the case that if there were no agents, I would not exist—nor would the cognitive structures in my mind/brain. But that’s not a reason for considering me or them problematic. It might be suggested that, for some things, it’s a necessary condition for their continued existence, not just that agents exist, but that agents *with specific kinds of purposes* exist. But that’s also a necessary condition of those purposes themselves existing. We cannot conclude that the purposes themselves are therefore problematic without undercutting what was supposed to be our basis for drawing the distinction between problematic and non-problematic forms of dependence.

Again, Jackendoff considers the Mississippi River and wonders where its boundaries are. Where, for example, do its tributaries end? Vagueness is indeed a complicated topic. But is Jackendoff suggesting that in any case where issues of vagueness arise it is suspect whether there are such objects? It would follow that the

existence of Ray Jackendoff is suspect. (Just what bits of matter are parts of Ray Jackendoff, and precisely when did he begin to exist?)³³ Issues concerning vagueness—as well as those concerning individuation, identity through time, relations to materially constituting parts, etc.—arise quite generally, including for objects as described in the sciences. Attempting to undermine IWS in this manner threatens to prove too much.

I conclude that Jackendoff has not established that IWS requires commitment to the existence of problematic objects.

3.2. *The notion of object*

Perhaps Jackendoff's objection fares better construed as a challenge to the IWS *notion* (or, concept) of object, according to which the very concept is unintelligible or unfruitful or in some other defective? It does not. One way he enters this charge is by denying that the IWS notion of object accords with a *commonsense* notion of object available to "direct intuition" (303) as perhaps he believes the notion of perceivable object is. We have already noted that it's unclear why IWS must limit itself to the notion of perceivable object and or to a notion available to commonsense. It could be that Jackendoff has an interest in undermining the *obviousness or inevitability* of IWS as a conception of semantics, so as to soften up opponents to CS as an alternative. But the original claim was that IWS is *incompatible* with CS, which is a different matter. Further, it's unclear in

³³ Cf. Unger (1979) and more generally on vagueness Keefe and Smith (1997). Jackendoff contrasts his problematic cases with those suggested by "an entirely intuitive construal" (300) of 'object.' But the objects on his list—Bertrand Russell, Noam Chomsky, trees, refrigerators, apartment buildings, screwdrivers—can be subject to boundary questions as well. Perhaps, in some of these cases, Jackendoff could try arguing that a "natural truth about the world," not some "matter of convenience," fixes the boundary—though I would not be optimistic. But this *clearly* would not work for artifacts like those on his list. Jackendoff later (351-2) brings up vagueness as a problem specifically for "necessary and sufficient conditions" analyses of concepts. But I think it fair to say that it presents a problem for everyone and that there's currently no consensus solution.

any event that the logical concept of object is *not* in some sense available to commonsense. This depends on what's meant by availability to commonsense. One might argue that individuals' mastery of the logical concept of object is displayed in their competence with quantificational and referential expressions and concepts, including their role in inference—a competence that arguably *is* available to a significant extent to reflective commonsense. Finally, we must distinguish a concept's intuitiveness from the intuitiveness of certain applications of it. That inquiry might yield some “curious beasts” (303) does not show that the *notion* of object is curious, just that some objects are. If Jackendoff wants to challenge the very notion of object employed by IWS, a list of intriguing cases will not suffice.³⁴

3.3. *Further remark on semantics and ontology*

Whatever its exact point, I do not believe Jackendoff's “objects” argument establishes that IWS cannot be reconciled with mentalism. However, I am sympathetic to the claim that a mentalist approach to semantics should not be held hostage to the fortunes of metaphysics. How we unreflectively conceptualize the world and how the world is need not align. Nor need either align with our *reflective* beliefs concerning how things are—in particular, with the sometimes counter-intuitive conceptions arrived at through empirical investigation. (I am also sympathetic to the claim, however, that there are limits, both empirical and conceptual, to the *extent* to which there can be

³⁴ There are indeed other arguments one might try. Gross (forthcoming-c) discusses Hilary Putnam's claim that the notion of object and cognate notions deployed in philosopher's ontological disputes are unintelligible owing to pervasive context-sensitivity. See, for example, Putnam (1999: 7-9). Chomsky's cases (2000: 126) of anaphors that seemingly shift reference in subtle ways ('He read the book and then burned it,' 'He painted the door and then walked through it') certainly pose challenges. There are at least first blush strategies in reply. For example, it's arguably part of one's knowledge of *books* that to burn a book-type is to burn one of its tokens. But Chomsky is certainly right to emphasize the complications.

misalignment.) A methodology that allows arguments from metaphysics into the mix when evaluating hypotheses concerning the psychological bases of semantic competence requires support.³⁵ Reference to distances may or may not be only apparent, but we should be wary of reaching a conclusion on the basis of antecedent ontological proclivities. Why suppose that the relevant psychological facts should be answerable to disputes over whether distances should be admitted into our ontology, even allowing the intelligibility of such disputes?

My sympathy stems, however, not from a Jackendoffian distrust of the apparent metaphysics of IWS, but rather—as my first reply to his argument underscored—from a conception of IWS that deflates its metaphysical pretensions. I do not question either the IWS notion of object or the objects that, according to IWS, speakers refer to. Rather, I suggest that IWS, rightly construed, can to a significant degree remain *agnostic* as to what *in fact* there is. There may well be as well a version of IWS with deeper metaphysical aspirations. I would then question its *relevance* to CS, not its compatibility with it.

I thus avoid the contentious metaphysical conclusion to which Jackendoff seems to think he is driven by the form his second objection to IWS takes—namely, that, having “abandon[ed] the unexamined notion of ‘objects in the world,’” we must “[push] ‘the world’ down into the mind.” (303) As Jackendoff recognizes (305), it’s difficult not to read these slogans as proclaiming a form of idealism—especially when he tells us that “the reality in which you are reading this book . . . is a product of our human modes of

³⁵ A particularly explicit example of such a methodology is Ludlow (1999: 5): “our task is to solve a kind of complex equation involving information from semantics on the one side and metaphysics on the other.” He winds up arguing for a treatment of tense and a metaphysics on which only the present is real. For some discussion, see Gross (forthcoming-a).

perception and conception.” (309)³⁶ But for CS to examine how humans conceptualize the world does not require in this sense abandoning objects or the world, or pushing them anywhere. For all Jackendoff has established, he too can for the most part remain agnostic, for his purposes, about whether and to what extent the conceptual structures he posits *in fact* accord with how things are.

But only for the most part, because of course he also holds, for instance, that the mind is the brain functionally characterized, and that the brain is a spatially extended object in a world that contains other things with which it interacts, including other brains (330). These are ontological commitments that Jackendoff readily accepts and that are central to his theorizing. He thus cannot “push” all of these other things “into the mind” without contradiction. In addition, he must allow that at least in *this* case how we conceptualize, experience, and talk about the world does indeed align with how it is.

It is *possible* that Jackendoff would relinquish his idealistic tendencies and claim that all he is *really* committed to are naturalistically acceptable objects: sure, *brains* exist (for emphasis, I suppose one might add: in the world), but they are part of our scientific world-view. I have already critically discussed Jackendoff’s apparent *physicalist* conception of naturalism. But we can now add that naturalism is in fact in *tension* with at least some of the arguments Jackendoff advances in support of his second objection to IWS. For the sciences speak of objects that are imperceptible and that are abstract; and

³⁶ One could agree if Jackendoff only meant that it’s a *partial* product of our human modes of perception and conception. After all, without these modes, there would not be books—at least none written by us. Books (*that* part of reality) require thinkers to write them. Maybe this is indeed all he means. Maybe, in fact, contrary to what is sometimes suggested, his talk of pushing things into the mind is not intended to entail anything at all about whether they exist. It might just be a colorful manner of turning attention away from objects and towards our concepts of them. I find his position unclear.

they employ vague terms so long as they are precise enough for our explanatory purposes. (What *are* the *exact* boundaries of the brain?)

4. Conclusion

Jackendoff wisely remarks that “in practice one cannot first establish the foundations and then go on to do the work. Rather, the empirical results are part of what motivates the search for new foundations The relation between the philosophy and the dirty work has to be a two-way street.” (268) In this paper, I have in effect argued that despite his good intentions, Jackendoff’s critique of IWS and its compatibility with mentalism is animated by dispensable philosophical assumptions: a physicalistic conception of naturalism, and a set of ontological assumptions both unmotivated in themselves and orthogonal to the goal of IWS as I see it. Making this explicit leaves us better placed to assess IWS as we ought: on its empirical merits—or demerits, as the case may be.³⁷

References

- Block, Ned (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy* 10: 615-678.
- Block, Ned (1990). Consciousness and accessibility. *Behavioral and Brain Sciences* 13: 596-7.
- Burge, Tyler (2003). Psychology and the environment: reply to Chomsky. In *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, Martin Hahn and Bjørn Ramberg (eds.), 451-70. Cambridge: MIT Press.

³⁷ [Acknowledgments: editor, anonymous referees, John Collins, Michael Weisberg; others]

- Burgess, John, and Gideon Rosen (1997). *A Subject with No Object: Strategies for Nominalistic Interpretation of Mathematics*. Oxford: Oxford University Press.
- Chomsky, Noam (1975). *Reflections on Language*. New York: Pantheon.
- Chomsky, Noam (1986). *Knowledge of Language*. New York: Praeger.
- Chomsky, Noam (2000). *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, Noam (2003a). Reply to Poland. In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 263-5. Oxford: Blackwell.
- Chomsky, Noam (2003b). Reply to Rey. In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 274-87. Oxford: Blackwell.
- Chomsky, Noam (2003c). Reply to Ludlow. In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 287-95. Oxford: Blackwell.
- Churchland, Paul (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy* 78: 67-90.
- Churchland, Paul (1989). On the nature of theories: a neurocomputational perspective. *Minnesota Studies in the Philosophy of Science* 14: 59-101.
- Crane, Timothy, and Hugh Mellor (1990). There is no question of physicalism. *Mind* 99: 185-206.
- Dennett, Daniel (1971). Intentional systems. *Journal of Philosophy* 68: 87-106.
- Fodor, Jerry (1974). Special sciences. *Synthese* 28: 97-115.
- Fodor, Jerry (1990a). A theory of content, II: the theory. In his *A Theory of Content and Other Essays*. Cambridge: MIT Press.

- Fodor, Jerry (1990b). Information and representation. In *Information, Language, and Cognition*, Philip Hanson (ed.), 175-91. Oxford: Oxford University Press.
- Fodor, Jerry (1998a). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Fodor, Jerry (1998b). Special sciences: still autonomous after all these years. *Philosophical Perspectives* 11: 149-63.
- Friedman, Michael (1974). Explanation and scientific understanding. *Journal of Philosophy* 71: 5-19.
- Gross, Steven (forthcoming-a). Can empirical theories of semantic competence really help limn the structure of reality? *Nous*.
- Gross, Steven (forthcoming-b). Context-sensitive truth-theoretic accounts of semantic competence. *Mind and Language*.
- Gross, Steven (forthcoming-c). Putnam, context, and ontology. *Canadian Journal of Philosophy*.
- Higginbotham, James (2004). Jackendoff's conceptualism. *Behavioral and Brain Sciences* 26: 680-1.
- Jackendoff, Ray (1991). The problem of reality. *Nous* 25: 411-34.
- Jackendoff, Ray (2002). *Foundations of Language*. Oxford: Oxford University Press.
- Jackendoff, Ray (2004). Toward better mutual understanding. *Behavioral and Brain Sciences* 26: 695-702.
- Kaplan, David (1989). Afterthoughts. In *Themes from Kaplan*, Joseph Almog, John Perry, and Howard Wettstein (eds.), 565-614. Oxford: Oxford University Press.
- Keefe, Rosanna, and Peter Smith (eds.) (1997). *Vagueness*. Cambridge: MIT Press.

- Kitcher, Philip (1981). Explanatory unification. *Philosophy of Science* 48: 507-31.
- Kripke, Saul (1982). *Wittgenstein on Rules and Private Language*. Cambridge: Harvard University Press.
- Larson, Richard, and Gabriel Segal (1995). *Knowledge of Meaning*. Cambridge: MIT Press.
- Loewer, Barry (1996). A guide to naturalizing semantics. In *The Blackwell Companion to the Philosophy of Language*, Bob Hale and Crispin Wright (eds.), 108-26. Oxford: Blackwell.
- Ludlow, Peter (1999). *Semantics, Tense, and Time*. Cambridge: MIT Press.
- Ludlow, Peter (2003). Referential semantics for I-languages? In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 140-61. Oxford: Blackwell.
- Millikan, Ruth (1989). Biosemantics. *Journal of Philosophy* 86: 281-97.
- Parsons, Charles (1982). Objects and logic. *The Monist* 65: 491-516.
- Partee, Barbara (1979). Semantics: mathematics or psychology? In *Semantics from Different Points of View*, Egli Bäuerle and Arnim von Stechow (eds.), 1-14. Berlin: Springer-Verlag.
- Pietroski, Paul (2003). The character of natural language semantics. In *Epistemology of Language*, Alex Barber (ed.), 217-56. Oxford: Oxford University Press.
- Poland, Jeffrey (2003). Chomsky's challenge to physicalism. In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 29-48. Oxford: Blackwell.
- Putnam, Hilary (1967). Psychological predicates. In *Art, Mind, and Religion*, William Capitan and Daniel Merrill (eds.), 37-48. Pittsburgh: University of Pittsburgh Press. Variousy reprinted as: The nature of mental states.

- Putnam, Hilary (1999). *The Threefold Cord: Mind, Body, and World*. New York: Columbia University Press.
- Quine, Willard van Orman (1960). *Word and Object*. Cambridge: MIT Press.
- Rey, Georges (1997). *Contemporary Philosophy of Mind*. Oxford: Blackwell.
- Rey, Georges (2003a). Chomsky, intentionality, and a CRTT. In *Chomsky and his Critics*, Louise Antony and Norbert Hornstein (eds.), 105-39. Oxford: Blackwell.
- Rey, Georges (2003b). Intentional content and a Chomskian linguistics. In *Epistemology of Language*, Alex Barber (ed.), 140-86. Oxford: Oxford University Press.
- Searle, John (1990). Consciousness, explanatory inversion, and cognitive science. *Behavioral and Brain Sciences* 13: 585-95.
- Soames, Scott (1985). Semantics and psychology. In *The Philosophy of Linguistics*, Jerrold Katz (ed.), 204-26. Oxford: Oxford University Press.
- Stalnaker, Robert (1997). Reference and necessity. In *The Blackwell Companion to the Philosophy of Language*, Bob Hale and Crispin Wright (eds.), 534-54. Oxford: Blackwell.
- Stich, Stephen (1983). *From Folk Psychology to Cognitive Science*. Cambridge: MIT Press.
- Stich, Stephen, and Ted Warfield (eds.) (1994). *Mental Representation*. Oxford: Blackwell.
- Stoljar, Daniel (2001). Physicalism. In *The Stanford Encyclopedia of Philosophy* (Spring 2001 Edition), Edward Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2001/entries/physicalism/>.

Unger, Peter (1979). I do not exist. In *Perception and Identity*, Graham Macdonald (ed.), 235-51. London: Macmillan.

Wright, Crispin (2002). What could antirealism about ordinary psychology possibly be? *The Philosophical Review* 111: 205-33.