

8

Psychological Reality of Grammars

Robert J. Matthews

1 INTRODUCTION

It is a measure of the depth of the conceptual revolution wrought by Noam Chomsky in linguistics that few linguists would quarrel with his notion that theoretical linguistics is a subfield of psychology.¹ Specifically, theoretical linguistics is taken to be that part of psychology concerned with human linguistic competence, viz., "the system of rules and principles that we assume have, in some manner, been internally represented by the person who knows a language and that enable the speaker, in principle, to understand an arbitrary sentence and to produce a sentence expressing his thought" (1980: 201). So construed, linguistics is not directly concerned with linguistic performance; however, it is assumed that models of language use will incorporate these knowledge representations:

A generative grammar is not a model for a speaker or a hearer. It attempts to characterize in the most neutral possible terms the knowledge of the language that provides the basis for actual use of language by a speaker-hearer. . . . No doubt, a reasonable model of language use will incorporate, as a basic component, the generative grammar that expresses the speaker-hearer's knowledge of the language; but this generative grammar does not, in itself, prescribe the character or functioning of a perceptual model or a model of speech production. (1965: 9)

Many philosophers, psychologists, and AI-researchers, and indeed more than a few linguists, remain deeply skeptical of Chomsky's linguistics-as-psychology program. These critics find his claim for the "psychological reality" of the grammars made available by linguistic theory empirically unsupported, if not simply unfounded, especially when one considers the empirical evidence upon which linguistic theory is apparently based.

In this paper I wish to defend Chomsky's claim. The argument that I sketch rests on considerations regarding the sort of relation of grammars to

performance models that would justify the claim that grammars are psychologically real. Grammars true of a speaker/hearer, I argue, bear such a relation to performance models. I begin by considering Chomsky's own defense of his claim. I then turn to Bresnan and Kaplan's (1982) argument against the psychological reality of transformational grammars. My criticism of their notion of psychological reality provides a context in which I can sketch my own argument. I conclude with reasons for thinking this argument is *largely* consistent with Chomsky's. What differences there may be between our respective arguments focus on the psychological reality not of the grammars but of the constructs that figure in those grammars.

2 CHOMSKY'S ARGUMENT FOR THE PSYCHOLOGICAL REALITY OF GRAMMARS

Chomsky's response to his critics (1980: 106ff) has been to deny what seems to be presumed by their arguments, namely, that there is a principled epistemological distinction between evidence that counts towards the psychological reality of the constructs of linguistic theory and evidence that counts "merely" towards the truth of that theory. Chomsky argues that there cannot be any such epistemological distinction, since the existence of that distinction would entail that there is a principled distinction between a linguistic theory's constructs being psychologically real and the theory's being true, and this second distinction he also denies (1980: 107).

Chomsky (1980: 189ff) contends that it makes no more sense to question the psychological reality of the theoretical constructs of linguistic theories that we accept as true than it does to question the physical reality of the theoretical constructs of physical theories that we accept as true. But is this true? Questions of the latter sort would presumably arise only in the context of a metaphysical discussion of the existential commitments of a physical theory that we accept as true, since in all other contexts, *physical* reality is what true *physical* theories are taken to describe. Questions of the former sort, however, could very well arise in non-philosophical contexts in which we are concerned *not* with the question of whether we should accept the existential commitments of our best linguistic theories (clearly we should, but see pp. 26ff), but with the *psychological import*, if any, of those commitments.

This, I take it, is the point of questioning the psychological reality of grammars, rules, representations, etc.: it does not follow from the fact that our best linguistic theory commits us to the existence of such *linguistic* entities that they have any *psychological import*, i.e. have any significance for the explanatory concerns that define psychology. Claims for the psychological reality of linguistic constructs will therefore be justified only insofar as we are able to establish their psychological import.

3 BRESNAN AND KAPLAN ON THE PSYCHOLOGICAL REALITY OF GRAMMARS

Bresnan and Kaplan's (1982) criticism of Chomsky's argument for the psychological reality of grammars rests on their assumption that generative-transformational grammars have not been successfully incorporated in "psychologically realistic" models of language use (1982: xvii). Bresnan and Kaplan, hereafter B&K, see Chomsky's argument as an attempt to preserve the claim that generative-transformational grammars represent a *psychological* hypothesis despite these grammars' failure to be incorporated in any realistic models of language use. Against Chomsky's claims for the psychological reality of transformational grammars, B&K's basic contention is this:

Linguistically motivated descriptions of a language need not bear any resemblance to the speaker's internal description of the language. Therefore, one cannot justifiably claim "psychological reality" for a grammar (in any interesting sense) merely because the grammar has some linguistic motivation. (1982: xxi-xxii)

According to B&K, justifiable claims for the psychological reality of a grammar require more than evidence that the grammar provides us with a description of the linguistic knowledge domain; they require evidence that "the grammar corresponds to the speaker's *internal* description of that domain" (1982: xxiii). More precisely, B&K require that psychologically real grammars be able to play the appropriate role in models of language use that satisfy what they call the "strong competence hypothesis," which they define as follows:

Suppose that we are given an information-processing model of language use that includes a processor and a component of stored linguistic knowledge K, [where] as a minimum . . . K prescribes certain operations that the processor is to perform on linguistic representations, such as manipulating phrases or assigning grammatical functions. . . . We call the subpart of K that prescribes representational operations the *representational basis* of the processing model. (The representational basis is the "internal grammar" of the model.) . . . A model satisfies the *strong competence hypothesis* if and only if its representational basis is isomorphic to the competence grammar. (1982: xxxi)

This construal of the notion of "psychological reality" provides the basis of B&K's argument against generative transformational theory. Schematically, their argument runs as follows:

- 1 it is an adequacy condition on linguistic theory that the grammars made available by a theory be psychologically real;
- 2 such grammars can be psychologically real only if they can satisfy the (strong) competence hypothesis; hence,

- 3 the grammars made available by an adequate linguistic theory must satisfy the competence hypothesis;
- 4 transformational grammars have not, and indeed cannot, satisfy the competence hypothesis (as is evidenced by the failure of the so-called Derivational Theory of Complexity); hence,
- 5 the grammars made available by transformational theories are not psychologically real, and those theories themselves fail a basic adequacy condition.

The crux of B&K's argument has to do with their claim that a grammar can be psychologically real only if it can satisfy their strong competence hypothesis. This is a strong claim indeed, for unlike Chomsky's notion of psychological reality which levies no specific performance-related requirements on psychologically real grammars, B&K's claim levies a very specific and hence restrictive requirement on the role that a psychologically real grammar must be able to play in a model of language use. In particular, to be psychologically real, a linguist's grammar must be isomorphic to the representational basis – the "internal grammar" – of a processing model of language use for that language.

B&K offer little by way of a justification of their strong competence hypothesis. They suggest (1982: xix) that to reject the hypothesis is "to adopt the theoretical alternative that a different body of knowledge of one's language is required for every type of verbal behavior," but why should this be so? Why couldn't there be a single body of linguistic knowledge that is recruited in the production of every type of verbal behavior, but that is not isomorphic to the grammar made available by the best linguistic theory? There would seem to be no reason whatever, unless B&K assume that an isomorph of this grammar will for other reasons have to be internally represented anyway, so that the issue here is simply one of parsimony of knowledge structures. This does indeed seem to be their assumption. They assume without argument that what is acquired in acquiring a language is an isomorph of the linguist's grammar. Given this assumption, it is not surprising that B&K would embrace their competence hypothesis, for they would certainly not wish to be driven to the view of Fodor, Bever, and Garrett (1974: 370ff) according to which language acquisition involves the internalization of a grammar of the sort made available by linguistic theory, yet that internalized grammar is never used in language processing. It would certainly be preferable to assume in the absence of contrary evidence that the same knowledge structure that is acquired in the course of learning a language is used in language processing.

4 MUST THE GRAMMARS OF LINGUISTICS BE INTERNALLY REPRESENTED?

B&K take it to be uncontroversial that "stored knowledge structures underlie all forms of verbal behavior" (1982: xix). Their assumption is *not*

uncontroversial;² but even if it were, it would not suffice to establish the premise that their argument for the competence hypothesis apparently needs, namely, that isomorphs of the grammars of linguistic theory must be internally represented. In the absence of an argument for this premise, B&K effectively lack any argument for taking the grammars of linguistic theory to be isomorphic representations of the knowledge structures – the internal grammars – that are said to figure in these processes.

Others have attempted to establish the needed premise. Fodor, Bever, and Garrett (1974), for example, argue that the existence of linguistic universals requires that an explicit representation of the grammar made available to linguistic theory be part of the language processor:

There are linguistic universals which serve precisely to constrain the form in which information is represented in grammars (i.e., the form of grammatical rules). The question is: If the universals do not also constrain the form in which linguistic information is represented in a sentence-processing system, how is their exercise to be explained? Surely, if universals are true of anything, it must be of some psychologically real representation of that language. But what could such a representation be if it is not part of a sentence encoding-decoding system? (1974: 369–70)

This line of argument suffers from an obvious defect: the conclusion depends on an equivocation between grammars as representations of linguistic knowledge (what the linguist constructs) and grammars as that knowledge itself (what's in the head). The argument fails to establish what is at issue, namely, that the speaker/hearer uses an explicit representation of his knowledge of language that is isomorphic to the linguist's grammar. In fact, the argument fails to establish that the speaker/hearer has as a component of his language processing mechanism *any* explicit representation of his grammatical knowledge. Consider a class of automata, each of which has been hardwired to recognize a particular member of a class of formal languages. A formal characterization of this class would mention certain linguistic universals true of each of the automata; yet none need contain an explicit representation of these universals or of the language they recognize, any more than a hand calculator must contain an explicit representation of the axioms of arithmetic.

A more subtle argument for the assumption that B&K need rests on a defense of what Jerry Fodor with characteristic impartiality calls the "Right View," which holds the following:

(a) Linguistic theories are descriptions of grammars. (b) . . . learning one's native language involves learning its grammar, so a theory of how grammars are learned is *de facto* a (partial) theory of how languages are learned. (c) . . . the grammar of a language is internally represented by speaker/hearers of that language; up to dialectical variants, the grammar of a language is what its speakers have in common by virtue of which they are speaker/hearers of the *same* language. (d) . . . the intentional representation of the grammar (or, equivalently for these purposes, the internally represented grammar) is

causally implemented in communication exchanges between speakers and hearers insofar as these exchanges are mediated by their use of the language that they share; talking and understanding the language normally involve exploiting the internally represented grammar. (1981a: 199)

The Right View clearly entails B&K's assumption and indeed their competence hypothesis itself, for as Fodor points out, the view "construes learning a language as a process that eventuates in the internal representation of a grammar, and it construes the production/perception of speech as causally mediated by the grammar that the speaker/hearer learns" (1981: 201).

But appellation notwithstanding, why should we believe the Right View to be the right view? In particular, why should one endorse the notion that the speaker/hearer's knowledge of language takes the form of an internally represented grammar of the sort made available by linguistic theory? Fodor's reply, if I understand him correctly, is that the assumption that grammars are internally represented is warranted by a Realist principle to the effect that one should accept the ontology that the best explanation presupposes. Specifically, the appropriate form of argument for the assumption that grammars are internally represented is to show that this assumption, when taken together with independently motivated theories of the character of other interacting variables (such as memory limitations and the like), yields the best explanation of the data about the organism's mental states and processes and/or the behaviors in which such processes eventuate.³

While I am inclined to endorse a Realist principle of the sort that Fodor invokes in defense of the Right View, I do not think that it can be used to support the claim that the language processor incorporates an internal representation of the grammar postulated by linguistic theory. We do, of course, have linguistic theories that correctly predict many of the linguistic intuitions of speaker/hearers. And certainly the best explanation of their predictive success is that these speaker/hearers have the linguistic *knowledge* that is represented by the grammars attributed to them. But nothing, so far as I can see, suggests that these grammars are internally represented by speaker/hearers, if by this one means explicitly represented or tokened. Given what little we know about the computational organization of the brain, we are simply not in a position to say how the linguistic knowledge represented by means of a grammar is realized and used computationally. Indeed, as Stabler (1983) has argued, given what we know about other, better understood computational systems, it seems reasonable to suspect that grammars are *not* explicitly represented at all: the limited plasticity (and lability) of acquired grammatical competence would seem to render explicit representation unnecessary, while the relatively greater efficiency of processors whose programs are "hardwired" rather than explicitly represented would seem to render it undesirable. But here again, these are speculations whose resolution must await further evidence regarding the computational organization of the brain. The present point is simply that

Realist scruples argue neither for the Right View nor for the competence hypothesis that this view entails. Of course, these scruples *do* dictate that we take the grammar postulated by an explanatorily adequate linguistic theory to have the speaker/hearer as a model, but this does not entail that this model incorporates an explicit representation of that grammar. The grammar could be realized in any way whatever, so long as the realization preserved the truth of the claim that the grammar represents the speaker/hearer's linguistic knowledge. The realization might be extremely abstract, by which I mean that there might be no answer to the question "What structures and/or processes of the language processor specifically represent the speaker/hearer's knowledge of language?"

To Representationalists, the suggestion that this question may have no answer seems tantamount to admitting that there may be nothing for the grammar made available by the true linguistic theory to be true of, except the behavior of the individual to whom the grammar is ascribed. Thus, for example, Fodor writes:

If, then, the notion of internal representation is *not* coherent, the only thing left for a linguistic theory to be true of is the linguist's observations (*de facto*, the intentions of the speaker/hearer as extrapolated by the formally simplest grammar). Take the notion of internal representation away from linguistic metatheory and you get positivism by subtraction. (1981a: 201)

Yet the options here are not only representationalism or postivism (more specifically, behaviorism). There is the option mentioned above: the grammar ascribed to an individual can be true of an individual, though not in virtue of any explicit representation in the individual of that grammar. Rather than ascribing any particular computational structure to the individual, the ascription of a grammar would simply ascribe to the individual the property of being able to recover and use the grammatical information marked by the structural descriptions that the speaker/hearer's grammar associates with sentences of his language. The ascribed capability is clearly not behavioral, since there is no presumption that an individual to whom this sort of capability is ascribed is able to manifest it.

5 TRANSPARENCY AS THE MEASURE OF PSYCHOLOGICAL REALITY

There are many different ways in which a speaker/hearer might realize a grammar (or, equivalently, the grammar be true of that speaker/hearer), none of which would have to involve the individual's having an explicit representation of that grammar. Some of these ways might involve the individual's having no "internal grammar" whatever; others would involve his having such a grammar, but not one isomorphic to the linguist's grammar that the individual realizes.

The Marcus (1980) parser is an example of the latter sort. That parser,

which implements a version of Chomsky's Extended Standard Theory (EST), incorporates a grammar, i.e. a rule system that governs the interpretive processes of the parser; however, it is *not* a grammar of the sort made available by EST (or any other linguistic theory, for that matter). The incorporated grammar, as Marcus explains, is a set of pattern-action rules similar to the rules of Newell and Simon's production systems: each rule is made up of a *pattern*, which is matched against some subset of contents of the input buffer and active node stack (the two data structures maintained by the interpreter), and an *action*, which is a sequence of operations on these contents. These pattern-action rules are quite different from the phrase-structure and transformational rules of EST. The most notable difference between the grammar incorporated in the Marcus parser and those made available by EST is that the former's pattern-action rules reflect in direct fashion assumptions about the design of the parser, namely, that it maintains two data structures, that only certain contents of these structures are accessible, and so on. The grammars of EST, by contrast, do not wear their algorithmic implementation on their sleeves – there is no commitment within the theory as to how the knowledge characterized by the grammar made available by this theory is implemented or used by speakers.

Although the Marcus parser does not incorporate (in any usual sense of the word) an EST grammar for English, EST does bear an *explanatorily transparent* relation to the parsing theory that would have this parser as one of its models. By this I mean that the syntactic generalizations that are captured by means of the theoretical constructs of EST (e.g. rules, principles, and structures) are *explained* in terms of the organization and operation of the mechanisms postulated by the parsing theory. These generalizations are explained in the straightforward sense that one can see, for example, that the generalizations stated in an EST grammar for English would be true of a speaker who incorporated a Marcus parser (or, to put it another way, EST is true of all models of the parsing theory); moreover, and more importantly, one can see *why* these generalizations would hold for such a speaker. Thus, for example, in the version of EST that Marcus's theory satisfies, passive constructions involve the application of a transformation rule ("Move NP") that moves a postverbal NP into subject position, leaving a phonetically unrealized trace in the postverbal position that is co-indexed with the moved NP. The Marcus parser builds the same EST-annotated surface structure, not by actually moving an NP from a postverbal position but rather by creating an appropriately co-indexed trace in the postverbal position after encountering a verb with passive morphology. If Marcus's theory of parsing were true of speaker/hearers, then that theory would provide a detailed explanation of why EST was true of them, too, since EST bears this explanatorily transparent relation to Marcus's theory.

B&K's (strong) competence hypothesis attempts to guarantee the psychological reality of grammars by requiring that the grammars made available by an adequate linguistic theory be such that the correct parsing

theory will bear an explanatorily transparent relation of a very particular sort to that linguistic theory. Specifically, the hypothesis requires that a psychologically real grammar satisfy the following condition: a model of the parsing theory – i.e., a parser for a particular language – must include a knowledge structure isomorphic to the grammar for the language parsed that as a minimum prescribes the operations that the processor/interpreter is to perform on linguistic representations. It is unclear why B&K would suppose that psychological reality would require that the operations of the parser be isomorphic to the derivations of the grammar, much less that these operations be controlled by a data structure isomorphic to the grammar, i.e. by an explicit representation of the grammar. Why shouldn't an explanatory transparency of the sort exhibited by EST suffice? After all, the various constructs that appear in an EST grammar receive an *interpretation* in the Marcus parser, in the straightforward sense that by appeal to the organization and operation of the parser we can explain why the EST grammar is true of a speaker/hearer who realizes the Marcus parser. Certainly, the rationale that B&K actually offer for their competence hypothesis is not compelling: there is, we have seen, no reason to suppose that the unification of linguistic knowledge structures requires it. The true rationale, I suspect, has more to do with *constraints* on linguistic theory: B&K want a methodological principle that will guarantee the pertinence of psycholinguistic experimental results, notably those involving measures of reaction time, to the problem of choosing between competing linguistic theories. This goal, they realize, requires the isomorphism that they postulate.

B&K's idea is not new. Miller and Chomsky (1963) proposed a similar competence hypothesis, arguing that if the isomorphism held, then the psychological "plausibility" of proposed grammars would be strengthened, since "our performance on tasks requiring an appreciation of the structure of transformed sentences [would be] some function of the nature, number, and complexity of the grammatical transformations involved" (1963: 481). If B&K's competence hypothesis were taken as a methodological principle, then experimental evidence regarding our performance on such tasks could be brought to bear on proposed grammatical theories in a straightforward way. In fact, the task of bringing such psycholinguistic evidence to bear on grammatical theories is *not* straightforward, even if the hypothesized isomorphism holds. Real time, as measured in reaction time experiments, need not bear any simple relation to "algorithmic time," as measured by the number of steps executed in the course of a computation. In order to bring psycholinguistic evidence to bear on the evaluation of grammatical theories, B&K's competence hypothesis would have to be supplemented with a theory of human computational complexity that would relate real time with algorithmic time. (Such a theory would minimally specify the computational architecture, the time and resource costs for primitive machine operations, and the implementation of the algorithm on that machine.) A theory of the requisite sort is presently beyond reach: we know very little about the computational machinery involved in language

processing. We are not therefore in a position to use experimental evidence regarding language processing, even if the isomorphism postulated by B&K's competence hypothesis obtains.

Berwick and Weinberg (1984) argue that it is unlikely that models of a plausible parsing theory of natural language will be isomorphic realizations of the grammars for the languages they parse. They do acknowledge that there is nothing known to date that would preclude a *type-homomorphic* realization of the sort once endorsed by Bresnan (1978: 3), where grammatical rules are mapped into parser operations in such a way that different rule types of the grammar are associated with different operation types of the parser.⁴ Yet if the intent of B&K's competence hypothesis is to insure that psycholinguistic evidence can be brought to bear on grammatical theories, then a homomorphic realization will not suffice, since under such a realization the derivation complexity of the grammar need not reflect the algorithmic complexity of the parser. B&K's competence hypothesis thus fails as a methodological principle not simply for want of a theory of human computational complexity that would enable us to apply the principle, but also because it seems doubtful that theories within its domain of application are even in principle capable of satisfying it. (This raises the obvious question: Why have a methodological principle that no theory can satisfy?)

As a criterion of psychological reality, B&K's competence hypothesis fares little better: it is unreasonably stringent and fails to provide for the comparative assessment of the psychological reality of the different grammars. Consider a parser that incorporates heuristic procedures of the sort hypothesized by Fodor, Bever, and Garrett (1974). Such a device would, for the sentences that it parses, recover the structural descriptions that the grammar for the language parsed associates with those sentences, yet it would do so in a way that bears no relation to the way in which derivations are constructed using the grammar. The generalizations stated in the linguistic theory of the grammar would be *satisfied* by models of this heuristic theory; however, they would not be explained by these models. Suppose this heuristic theory provided a correct account of language processing. The relation of parser to grammar would be explanatorily opaque, yet the grammar might nevertheless be regarded as psychologically real, inasmuch as the grammar correctly specifies the descriptions under which linguistic utterances are processed and interpreted. Fodor, Bever, and Garrett (1974) envisage just such a possibility in their discussion of the psychological reality of then proposed grammars. They argue that although there is little evidence for "the computational processes specified by transformational grammars," the structural descriptions specified by these grammars are psychologically real, since "the parameters that these descriptions mark enter, in one way or another, into a variety of psychological processing concerning language" (1974: 273-4).

Fodor, Bever, and Garrett's remarks suggest the usefulness of being able to distinguish both degrees and kinds of psychological reality, something that B&K's competence hypothesis fails to do. For reasons that will

become clear shortly, I propose to distinguish two kinds of psychological reality: that of grammars, and that of the linguistic constructs to which grammars advert. Truth, as Chomsky argues, is the relevant measure for the psychological reality of grammars. Explanatory transparency is a reasonably good measure for the psychological reality of the grammar's constructs, since the more explanatorily transparent the relation of parser to grammar, the more information (roughly speaking) the grammar carries regarding the structure of the parser.

6 CHOMSKY: GRAMMARS AS PSYCHOLOGICAL HYPOTHESES

Chomsky often writes as if he were a proponent of what Fodor calls the "Right View." He speaks of mental computations involving rules, of the mental representations that these computations provide, and of the role of these computations in the etiology of behavior. Thus, for example, in concluding an extended example in which our inability to form a particular question is explained in terms of the so-called *wh*-island constraint, Chomsky says the following:

Tentatively accepting this explanation, we impute existence to certain mental representations and to the mental computations that apply in a specific way to these mental representations. In particular, we impute existence to a representation in which (12) [_s which for PRO to play Sonatas on t]] appears as part of the structure underlying (5) [What sonatas are violins easy to play on?] at a particular stage of derivation, and to the mental computation that produces this derivation, and ultimately produces (5), identified now as ungrammatical because the computation violates the *wh*-island constraint when the rule of *wh*-movement applies to *sonatas* in (12). We attribute psychological reality to the postulated representations and mental computations. In short, we propose . . . that our theory is true. (1980: 196-7)

These conclusions certainly sound like an endorsement of the Right View, until, reading further, we learn that unlike proponents of the Right View who take themselves to be describing certain computational mechanisms and processes, Chomsky takes the linguist to be describing "abstract conditions that unknown mechanisms must meet" (1980: 197). When we couple this characterization of the linguist's project with Chomsky's often repeated insistence that a generative grammar is not a model for a speaker/hearer, that it does not prescribe the character or functioning of such a model, then we should begin to suspect that the so-called Right View is not Chomsky's view. But what, then, is his view; and how is his talk of computations and representations to be understood?

The mental computations to which Chomsky refers are, as he says in the above quotation, those that produce a particular syntactic *derivation*; the mental representations are the *structures* that appear in a derivation as a result of these computations. Chomsky ascribes psychological reality to these computations and representations, yet at the same time he insists that

the derivations in which these computations and representations figure do not pretend to model the psychological processes of a speaker/hearer: "When we say that a sentence has a certain derivation with respect to a particular generative grammar, we say nothing about how the speaker or hearer might proceed, in some practical or efficient way, to construct such a derivation" (1965: 9). The psychological reality that he ascribes to these constructs is not that bland, uncontentious reality to which he sometimes appeals when answering his critics, namely, that of these constructs being *true* (or *true-of-an-individual*). Rather, as Chomsky's criticism of Dummett makes clear, it is a psychological reality that would justify claims to the effect that linguistic theory is a *psychological* hypothesis.

Dummett, it will be recalled, claimed that his theory of meaning is not a psychological hypothesis because "it is not concerned to describe any inner psychological mechanisms" (1976: 70). Chomsky rejects this line of reasoning, arguing that "Dummett's theory of meaning is a 'psychological hypothesis,' though one that abstracts away from many questions that can be raised about inner mechanisms" (1980: 111). Dummett's theory is a psychological hypothesis, Chomsky argues, because it specifies *conditions* that the "inner psychological mechanisms" are alleged to meet.

Grammars are psychological hypotheses in precisely the same sense: they specify conditions that inner psychological mechanisms of the speaker/hearer are alleged to meet. A grammar does this, Chomsky argues, by specifying *intensionally* the function that these inner mechanisms are alleged to compute:

The grammar is a system of rules and principles that determines a pairing of sound and meaning (or better, a pairing of conditions on sound and meaning given by appropriate representations). The grammar is a function in intension, though this remark is misleading, since the grammar has many important properties beyond specifying a language in extension and in fact may not even specify such a language. (1980: 82)

The notion that a grammar is an intensional specification of a function is, as Chomsky says, misleading if we conclude from this that the grammar is merely a way of providing a finite specification of antecedently given language: "Since the language is infinite, it makes no sense to speak of it as 'given', except insofar as some finite characterization – a function in intension – is given" (1980: 84). The point here, I take it, is that what is given are speaker/hearers, not languages (construed as infinite sets), and grammars provide a characterization of these speaker/hearers by specifying the function (i.e. the pairing of sound and meaning) that they compute in the course of language use.

On what I take to be Chomsky's intended construal of grammars as intensional specifications of the function that speaker/hearers compute (hereafter, the "intended construal of grammars"), he is clearly justified in construing grammars as psychological hypotheses, despite his unwillingness to interpret these grammars as making any performance claims.

Specification of the function computed in the exercise of a cognitive capacity is, as many have emphasized, a crucial step in the development of a psychological theory of that capacity; indeed, if Marr (1982) is correct, specification of the function computed – what he called the “theory of the computation” – is *the* crucial step in the development of such a theory, since in the absence of such a specification, the inquiry into psychological mechanisms is in most domains hopelessly unconstrained. Chomsky endorses claims by Marr (e.g., 1982: 28) that linguistic theory is such a theory of computation, i.e., that it is concerned with the function computed by speaker/hearers and not with the algorithms or mechanisms by which that function is computed. Citing Marr, he says:

We may consider the study of grammar and UG to be at the level of the theory of the computation. I don't see any useful distinction between “linguistics” and “psychology,” unless we choose to use the former term for the study of the theory of the computation in language, and the latter for the theory of the algorithm. (1980b: 48–9)

On the intended construal of grammars, Chomsky is surely also justified in insisting on the “psychological reality” of grammars. To the extent that a grammar correctly specifies the function computed by a speaker/hearer, the grammar is true of that speaker/hearer. What point could there be in denying the psychological reality of such a grammar, except on the grounds that it is false of the speaker/hearer, i.e., that it provides an incorrect specification of the function computed by the speaker/hearer? It is not as if we have any other way of specifying the function. Certainly we cannot specify the function extensionally, since the set that defines the function is infinite. In principle we could specify the function intensionally by specifying the mechanisms that compute the function. But surely we are not now in a position to provide such a specification, and even if we were, that would not impugn the psychological reality of the grammar in question, since the grammar would still be true of the speaker/hearer; it would still provide a correct specification of the function computed by the speaker/hearer, albeit at a more abstract level of description.

Chomsky's construal of transformational grammars as specifications of the functions computed by a speaker/hearers has not gone unchallenged. Stabler (1984) argues that the construal is untenable. He points out that the domains of the grammars made available by current transformational theory are not coextensive with the domains of the parsing mechanisms that speaker/hearers embody; indeed, their domains fail to be coextensive even under idealizations of the parsing mechanisms that abstract away from various performance limitations (e.g. memory limitations). Thus, speaker/hearers both succeed in understanding certain ungrammatical sentences that fall outside the domain of the grammar attributed to them and at the same time fail to understand certain grammatical sentences that fall within the domain of that grammar. Moreover, ambiguous sentences get multiple representations within the grammar, while empirical evidence

seems to indicate that speaker/hearers typically compute only a single representation.

By itself, the apparent disparity of domains does not impugn Chomsky's construal of transformational grammars, since it might simply be attributed to errors in current linguistic theory. Chomsky, after all, is not committed to the view that the grammars made available by a false linguistic theory specify the functions computed by speaker/hearers. Stabler attempts to buttress his argument against this line of response by arguing that the disparity manifests methodological commitments in transformational linguistics which are inconsistent with the computational-theoretic enterprise described by Marr. The methodological commitments that he discusses, most notably the emphasis upon the formal simplicity of proposed theories, may very well impugn the attempts of Berwick and Weinberg (1984) to offer a computational-theoretic interpretation of certain of the linguistic constructs to which transformational grammars advert (viz. subjacency); however, such commitments will serve to impugn Chomsky's construal of transformational grammars only if it can be demonstrated that they lead to the specification of functions that speaker/hearers are unable to compute efficiently over the range of sentences that these speaker/hearers can use and understand. In the absence of such a demonstration, Stabler would seem to have no grounds for impugning Chomsky's construal. Nor does it seem likely that such a demonstration will soon be forthcoming. Given our impoverished understanding of human computational architecture, we are not now in a position to provide a demonstration based upon empirical considerations regarding human computational powers. And for the reasons detailed by Berwick and Weinberg (1984), a compelling demonstration based upon mathematical parsability results seems extremely unlikely; the application of those results to human language processing is tenuous in the extreme.

7 PSYCHOLOGICAL REALITY OF LINGUISTIC CONSTRUCTS

While on their intended construal, grammars true of a speaker/hearer are psychologically real, it would seem an open question whether the linguistic constructs to which grammars advert, including rules, representations, and the computations that figure in syntactic derivations, are psychologically real. Chomsky (1980: 197) argues that we are justified in attributing psychological reality to the constructs postulated by a grammar true of speaker/hearer. In effect, the psychological reality of these constructs is assumed to be inherited from that of the grammar. But this assumption seems arguable. It might be objected that in the absence of specific evidence for their existence, these constructs must be assumed to be artifacts of the particular way in which the function computed by the speaker/hearer is specified, and as such cannot inherit the psychological reality of the grammar in which they figure.

Realist principles, which I endorse, dictate that we should accept the

existential commitments of our best theories. Determining the existential commitments of our theories, however, is not the simple task that many assume. It is not just a matter of blindly Ramsifying our theories in the way that the slogan "to be is to be value of a bound variable" might *seem* to suggest. Suppose, for example, that the best "theory" of my present location on the Earth's surface includes a statement to the effect that I am presently located at 40 degrees 30.25 minutes North latitude, 74 degrees 26.04 minutes West longitude. Whatever the existential commitments of that theory, it is surely *not* committed to the existence of a certain quantity of something called "latitude" or "longitude." What the theory is presumably committed to is the existence of the physical location that is *specified* in terms of latitude and longitude. Latitude and longitude are *representational constructs* that figure in the system of representation that we employ in the theory to specify location. The point here is a very general one: in determining the existential commitments of a theory, we must distinguish the *theoretical magnitudes* to which the theory is existentially committed from the representational constructs to which the theory is *not* existentially committed and which serve only to specify the theoretical magnitudes. In most cases, the distinction between theoretical magnitudes and representational constructs is reasonably easy to draw: there is an intended interpretation of the theory, and that interpretation draws the required distinction. In practice, the distinction is often facilitated by the conventional nature of the system of representation in which the representational constructs figure: we recognize that the system could, in principle at least, be replaced by an alternative system of representation with different representational constructs without compromising the intended interpretation of the theory. (My present location, for example, could have been specified in polar coordinates or simply by street address.)

The distinction that the intended interpretation of a theory draws between theoretical magnitudes and representational constructs can be, and often is, redrawn by reinterpreting the status of certain terms in the theory. Typically this occurs as a consequence of new empirical and/or theoretical considerations that lead theorists to reconstrue the status of the theory's representational constructs. In the case of Balmer formula for the spectral lines of hydrogen, for example, the integers of that formula initially received no physical interpretation; however, when Bohr succeeded in deriving the generalized version of the Balmer formula from his theory of the atom, these integers were then interpreted as quantum numbers. Even more striking perhaps is the case of Maxwell's equations, in which the scalar and vector potentials were initially interpreted by virtually all physicists as representational constructs, but under the pressure of subsequent theoretical developments have come to be regarded as the essential existential commitments of the equations. Similar cases within the domain of linguistic theory are difficult to come by, for the simple reason that there are no theories or models of sentence-processing mechanisms that are sufficiently compelling to force an interpretation of the theoretical constructs that figure in linguistic theory.⁵

On their intended construal, therefore, grammars are not assumed to bear an explanatorily transparent relation to underlying performance mechanisms. As Chomsky himself puts it,

Although we may describe the grammar G as a system of processes and rules that apply in a certain order to relate sound and meaning, we are not entitled to take this as a description of the successive acts of a performance model. (1968: 117)

Chomsky's point, I take it, is that it is not part of the intended interpretation of grammatical theories that the linguistic constructs that figure in a grammar are to be construed as descriptions of underlying mechanisms. The intended interpretation is to take the grammar as simply the intensional specification of the function – "the pairing of sound and meaning" – that the speaker/hearer computes.

The conclusion that I wish to draw here should be clear: on the intended construal, what is taken to be psychologically real is not the linguistic constructs that figure in the grammar, but the function that these constructs serve to specify. These constructs, which include rules, representations, and the computations that figure in syntactic derivations, constitute the representational constructs of grammatical theory, and as such are *not* among the existential commitments of that theory. The function specified in terms of these constructs, on the other hand, *is* among the theory's existential commitments, and hence presumed to be psychologically real. Of course, evidence might be forthcoming which would force a reconstrual of certain, perhaps all linguistic constructs in the grammar. Specifically, we might obtain evidence establishing the existence of an explanatorily transparent relation between the grammar, on the one hand, and the computational mechanisms by which a speaker/hearer computes the function specified by the grammar, on the other. In such event, we would surely acknowledge the psychological reality of the linguistic constructs that receive an interpretation under the transparency relation. But the point to be emphasized here is that acknowledging the psychological reality of these constructs would involve a *reconstrual* of the grammar; it is not now part of the intended construal of grammatical theories that these constructs are presumed to be psychologically real. Whether this presumption will be overturned remains an open empirical question.

NOTES

- 1 By "psychology" we understand that scientific inquiry concerned with human cognitive capacities, with the mental structures that underlie those capacities, and with their exercise.
- 2 See Dennett (1983), Matthews (1984), Matthews (1988), and Stabler (1983).
- 3 I am here paraphrasing Fodor's (1981b: 120) account of what would constitute evidence for the claim that the postulates of logic are internally represented.
- 4 Bresnan's adequacy condition for "realistic" grammars:

We should be able to define for [a grammar] explicit realization mappings to psychological models of language use. These realizations should map distinct grammatical rules and units into distinct processing operations and informational units in such a way that different rule types of the grammar are associated with different processing functions. If distinct grammatical rules were not distinguished in a psychological model under some realization mapping, . . . the grammar could not be said to represent the knowledge of a language user in any psychologically interesting sense. (1978: 3)

- 5 Marcus (1980) and Berwick and Weinberg (1983, 1984) have claimed that a deterministic model of sentence processing explains the existence of, and some of the properties of the subadjacency constraint on natural languages imposed by current transformational theories; however, Fodor (1985) argues that the empirical arguments offered in support of this claim are flawed and the purported explanatory relationship between determinism and subadjacency is weak.

REFERENCES

- Berwick, R. and Weinberg, A. 1983: Syntactic constraints and efficient parsability. *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, 119-22.
- 1984: *The Grammatical Basis of Linguistic Performance*. Cambridge, Mass.: MIT Press.
- Bresnan, J. 1978: A realistic transformation grammar. In M. Halle, J. Bresnan, and G. Miller (eds), *Linguistic Theory and Psychological Reality*. Cambridge, Mass.: MIT Press, 1-59.
- Bresnan, J. and Kaplan, R. 1982: Introduction: grammars as mental representations of language. In J. Bresnan (ed.), *The Mental Representation of Language*. Cambridge, Mass.: MIT Press, xvii-iii.
- Chomsky, N. 1965: *Aspects of the Theory of Syntax*. Cambridge, Mass.: MIT Press.
- 1968: *Language and Mind*. New York: Harcourt, Brace, Jovanovich.
- 1980: *Rules and Representations*. New York: Columbia University Press.
- Dennett, D. 1983: Styles of mental representation. *Proceedings of Aristotelian Society*, 213-26.
- Dummett, M. 1976: What is a theory of meaning? (II) In G. Evans and J. McDowell (eds), *Truth and Meaning*. London: Oxford University Press, 67-137.
- Fodor, J. 1981a: Some notes on what linguistics is about. In N. Block (ed.), *Readings in the Philosophy of Psychology*, Vol. II. Cambridge, Mass.: Harvard University Press, 197-207.
- 1981b: Three cheers for propositional attitudes. In *Representations*. Cambridge, Mass.: MIT Press, 100-123.
- Fodor, J., Bever, T. and Garrett, M. 1974: *The Psychology of Language*. New York: McGraw-Hill.
- Fodor, J. D. 1985: Deterministic parsing and subadjacency, *Language and Cognitive Processes*, 1: 3-42.
- Marcus, M. 1980: *A Theory of Syntactic Recognition for Natural Language*. Cambridge, Mass.: MIT Press.
- Marr, D. 1982: *Vision*. San Francisco: W. H. Freeman.
- Matthews, R. 1984: Troubles with representationalism. *Social Research*, 51: 1065-97.

- (1988): The alleged evidence for representationalism. In S. Silvers (ed.), *ReRepresentation*. Dordrecht: Reidel, 103-72.
- Miller, G. and Chomsky, N. 1963: Finitary models of language users. R. Luce, R. Bush, and E. Galanter (eds.), *Handbook of Mathematical Psychology*, vol. II. New York: Wiley, 185-94.
- Stabler, E. 1983: How Are Grammars Represented? *The Behavioral and Brain Sciences*, 6: 391-402.