

KNOWLEDGE AND SEMANTIC COMPETENCE

I INTRODUCTION: WHAT IS THE PURPOSE OF A THEORY OF MEANING?

1.0 Motivation and General Considerations

This discussion is about linguistic competence – the ability of speakers to understand their language. Our focus, in particular, is on *semantic* competence, an ability to *interpret* language. To see its theoretical interest, consider an unusual description of a familiar type of phenomenon. John sees Mary searching for something in her living room. He surmises she has misplaced her scarf. Remembering recently having seen it under the table, he believes that if she knew what he remembered it would facilitate her search. He takes a short breath; the air in his lungs releases at a slow steady rate; his vocal folds contract and relax in an elaborate fashion; and as the air passes into his mouth, his jaw, lips and tongue move in complicated ways, all of which serve to create a specific vibratory pattern, which sounds like an utterance of, ‘I saw your scarf under the table’. The sound pattern bounces off sensitive bits of tissue in Mary’s inner ear, and shortly afterwards, her search ceases with the scarf recovered.

Many of these details are of theoretical interest, but our focus will be on what enables Mary to recognize that John’s utterance means at some time prior to it, *John saw Mary’s scarf under the table*. The most common answer is that linguistic competence equals *knowledge* of a theory, or ‘grammar’ (see §1.1). In what follows we will review influential answers to the question, ‘Are speakers able to understand their (first) language in virtue of bearing a doxastic relation to a grammar?’ Before we start, we will outline some technical terms we will employ throughout.

1.1 Terminology.

A *grammar* is an abstract entity; in particular, it produces syntactic structures of a language and assigns them meanings and phonological forms. Though only part of a grammar generates interpretations of sentences, we will frequently speak of grammars and knowledge of grammars. Occasionally, when more specificity is required, we will speak of semantics and semantic knowledge. A *psychogrammar* is a mental state of knowing a grammar (if such a state exists); it is ‘a mental condition on a par with the state of *thinking of the number 3*’ (George 1989b, 90). In short, a speaker’s grammar is an object he knows, and his psychogrammar is his state of knowing it. Thus, as George notes, ‘[w]e might come to be able to articulate the object of a speaker’s knowledge, the grammar, without thereby being able to say

how that object is represented by the speaker. The grammar is what is represented, not what is doing the representing' (George 1989b, 91).

A *physiogrammar* is a physical state (if such a state exists) of the speaker that realizes the psychogrammar. Just as a correct theory of a speaker's grammar does not render one theory of her psychogrammar more plausible than all others, a correct theory of her psychogrammar does not render one theory of her physiogrammar more plausible than all others either. On this picture, if knowledge of language enters into an explanation of behavior, then, if such explanations are causal, a psychogrammar enters into the explanatory causal chain. If the psychogrammar is identical to the physiogrammar, then of course the latter is a part of that causal chain. Grammars, though, since abstract, *cannot* be causally efficacious; they are objects of knowledge, and so they can be no more causally responsible for behavior than Santa Claus should little Billie become joyful when he anticipates that Santa Claus is coming to town (George 1989b, 92).

It is common ground that there is a systematic relationship between knowledge of a grammar (i.e., one's psychogrammar) and whatever other beliefs one forms as a result of linguistic competence. Suppose John hears Mary utter, 'Hesperus burns bright tonight', and according to his knowledge of grammar, (roughly) an utterance of 'Hesperus burns bright tonight' is true iff Hesperus burns bright on the evening of the utterance. He will, *ceteris paribus*, believe that Mary said Hesperus burns bright that night. A *processing algorithm* (if one exists) is an abstract object that describes processes of linguistic perception and production. Such an algorithm takes John from his grammar-induced belief and his perceptual belief (something to the effect that Mary produced an utterance of a certain form) to his belief about what was said.

In §2 we will discuss the plausibility of supposing speakers have some sort of knowledge of a grammar of their language. The views we consider address whether speakers are doxastically related to grammars of their languages, and if so, what the nature of that relation is. In §3, we focus on an argument designed to show that linguistic competence cannot be adequately explained by describing a grammar (as characterized above) and the relation one bears to it.

II TACIT KNOWLEDGE

2.0 Introduction: Cognitivism.

Do competent speakers know a grammar of their language? We speak of 'knowledge of meaning' and 'knowledge of language'. Yet whatever we mean by such locutions talk of this sort of knowledge differs from the knowledge that $2 + 2 = 4$, or that one's favorite cup is filled with coffee. Unlike the latter two, the former seems to be rarely (if ever) explicitly statable by its knower. Most speakers cannot state principles which would explain the ungrammaticality of (2.1) and the grammaticality of (2.2).

(2.1) *John believed that any senators were drunk.

(2.2) John doubted that any senators were drunk.¹

Most speakers, though capable of using (2.3) and (2.4), cannot explain why only (2.3) permits substitution of identities *salva veritate*.

- (2.3) Mary saw the student leave.
 (2.4) Mary saw that the student left.

If the student who left is the happiest girl in Newark, then Mary saw the happiest girl in Newark leave; not so for (2.4). Furthermore, unlike (2.3), (2.4) can be true even if Mary saw no one leave. Perhaps she noticed that the previously occupied chair was empty (for further discussion of differences between (2.3) and (2.4), see Higginbotham 1983; 1995).

Many authors dismiss the ascription of knowledge to speakers as 'unnatural' or 'incoherent', since we lack conscious access to it (Foster 1975, 2; cf. Schiffer 1987, 255-261, Dummett 1975). A traditional and still common way to deal with this problem is to attribute *tacit knowledge* of a grammar (Chomsky 1965, 8; 1986, 266). Speakers have propositional knowledge of a grammar, but such knowledge is inaccessible to consciousness. Speakers understand a sentence of their language, because they exploit a grammar to (unconsciously) compute a meaning theorem for the sentence. Positing tacit knowledge is justified if so doing explains linguistic behavior better than any rival account.

Any theory which treats speakers as linguistically competent *in virtue of* tacit knowledge of a grammar we shall call **cognitivism**. Cognitivism is the received view in linguistics, as can be seen by a glance at the introductory chapters to linguistics textbooks: (Culicover 1997, 1-3; Cowper 1992, 1-4; Hagemann 1990, Larson and Segal 1995, 9-22). In this section, we will discuss an attempt to justify cognitivism, as well as some famous objections. (We will then discuss theories like cognitivism classified according to the relation they posit between a semantic theory and a speaker's mental state.)

Two points about cognitivism are relevant. First, noted by Higginbotham (1994), talk of knowledge of meaning can be misleading: 'knowledge of meaning is a *phenomenon*, not a hidden *explanandum* [sic]. A psychology for me that simply omitted to state that I knew the words 'snow is white' meant in my speech that snow is white would be in so far forth a false psychology' (p. 88; cf. Segal 1994, 115-116). However, at stake is not whether linguistic competence *per se* requires semantic knowledge. Our concern is whether such competence requires (tacit) knowledge of *all* of a semantic theory, which, *prima facie*, is not what a psychological theory should predict. Secondly, 'cognitivism' is ambiguous between requiring that a grammar only specify knowledge one has of the structure, meaning, and phonological properties of sentences (a position endorsed by Chomsky (1965), pp. 8-9 and Samuel Keyser²) or requiring that whatever procedures a grammar uses to derive appropriate meaning theorems are psychologically real as well, in the sense that they are 'mirrored' by a process in one's mind (i.e., in one's psychogrammar) (cf. Davies' discussion of his 'mirror constraint' (1981, 53-55; 1987, 446-447; Chomsky 1986, 263-273). The latter requires (something like) propositional knowledge of the axioms of a particular meaning theory, whereas the former is quiet about the nature of linguistic knowledge. Our discussion of cognitivism will focus

exclusively on the latter, though much of what we say here and in a later discussion of dispositionalism will apply, *mutatis mutandis*, to the former view.

2.1 Justifying Tacit Knowledge of a Grammar.

In articulating how cognitivism might be justified, we appealed to a 'best theory' principle. While defending this principle would require delving into more philosophy of science than we have space for, Fodor's (1968) defense is worth commenting on. Fodor attempts to justify a general principle for positing tacit knowledge, the crux of which is that one way to explain how a type of behavior might occur is by building a machine that simulates the behavior. His argument divides into three stages. First, he argues that a computer's programming language 'can be thought of as establishing a mapping of the physical states of a machine onto sentences of English such that the English sentence assigned to a given state expresses the instruction the machine is said to be executing when it is in that state' (p. 639). Second, if the programmed machine 'optimally' simulates an organism's behavior, then the machine exhibits a type of behavior (if and) only if the organism does, and for each type the machine can exhibit, the sequence of (computationally relevant) states of the machine resulting in that behavior can be mapped onto a sequence of English sentences, such that the latter constitutes a true etiology of the machine's output. Finally, he invokes a general principle of inductive inference, namely,

If D is a true description of the etiology of an event e, and if e' is an event numerically distinct from e but of the same kind, then it is reasonable to infer, *ceteris paribus*, that D is a true description of the etiology of e' (p. 639).

He concludes,

If X is something an organism knows how to do, but is unable to explain how to do, and if S is some sequence of operations, the specification of which would constitute an answer to the question 'How do you X?', and if an optimal simulation of the behavior of the organism X-s by running through the sequence of operations specified by S, then the organism *tacitly knows* the answer to the question, 'How do you X?', and S is a formulation of the organism's tacit knowledge (p. 638).

To be sure, Fodor's defense is schematic. Filling in details would involve resolving a number of issues, for instance, what counts as behavior. Since Chomsky's review of Skinner's *Verbal Behavior* (Chomsky 1959), it has been widely acknowledged that there is more to behavior than what behaviorism included. But, as is also well known, including more than overt physical behavior in an *explanandum* engenders other sorts of problems. Another question requiring an answer concerns how to construct a theory of event types in a principled way so that relevant human and machine behaviors get typed together. This problem also increases in complexity when the extension of 'behavior' is expanded. A third question concerns what counts as optimal simulation of behavior? Since there have been but a finite number of human behaviors, there are infinitely many different ways of producing those behaviors. For that matter, there are infinitely many different ways of producing reasonable infinite extensions of those behaviors. So, beyond extensional equivalence, we need additional criteria for what counts as optimal simulation. What these criteria are and what justifies them is well nigh tantamount to explaining what

makes for a good theory, or why one theory is to be accepted over another. (Further discussion of Fodor's argument is found in Graves et al. 1973.)

Wright notes that machine simulations of complex behavior do not always license intuitively plausible ascriptions of tacit knowledge. It is possible to write a program that simulates a homing pigeon's ability to find its way home from indefinitely many distant locations, but Wright contends that the bird lacks any sort of tacit knowledge of a homing theory that issues in homing theorems about where it should fly to next (Wright 1986a, 41-42, 1986b, 235-37). (This type of argument will receive detailed discussion in §2.2 and §2.3.)

Quine put forward a powerful and influential objection to positing tacit knowledge of grammar (Quine 1972). First, he observes that any finitely axiomatizable theory can be finitely axiomatized in infinitely many ways. So, if there is one finite grammar of a language, there are infinitely many. Furthermore, such grammars are *extensionally equivalent*; they all generate the same sentences, and assign them the same meanings and phonological forms. Thus, if English has a finite grammar, it has infinitely many. Second, Quine distinguishes two relations one might bear to a grammar. In his terminology, either it *fits* the linguistic behavior of competent speakers; or, if an adult learned, say, English (for the first time) by memorizing a particular grammar, then that grammar – unlike extensionally equivalent ones – *guides* his behavior. Positing tacit knowledge of a grammar presumably amounts to linguistic competence in virtue of speakers being guided by the grammar in some sense.³ Thus, according to cognitivism, one grammar is 'special' in the sense that it is the one used, i.e., it correctly describes the mental processing that underlies sentence comprehension in a way that its extensional equivalents do not. But then even where there is complete agreement about sentences of the target language, one grammar still must be singled out from its extensional equivalents. What justifies selecting one over another? As Quine puts it,

If it is to make any sense to say that a native was explicitly guided by one system of rules and not by another extensionally equivalent system, this sense must link up somehow with the native's dispositions to behave in observable ways in observable circumstances (Quine 1972, 444).

Thus, the task Quine sets for the cognitivist is to find 'a criterion of what to count as the real or proper grammar, as over against an extensionally equivalent counterfeit' (p. 448; cf. George 1986, 493-496 for further discussion of how the ascription of tacit knowledge of a grammar is not fully justified by the kind of behavioral data Quine is concerned with).⁴

As a point of scholarship, Quine's wording is ambiguous. His text supports characterizing the project as what Davies calls 'Quine's challenge', which involves answering how there *can* be empirical evidence to warrant attributing tacit knowledge of one theory rather than another, extensionally equivalent, one (Davies 1987, 442). But it also supports a reading under which Quine's attack on tacit knowledge centers around the plausibility of there actually being evidence favoring one grammar over its extensionally equivalent counterfeits.⁵ One could satisfy the former and not the latter. In some possible world, when supplied with hypnotic suggestion, we immediately write down a particular grammar. This scenario only shows how there can be empirical evidence, not that there is empirical evidence. The latter view seems more Quinean in spirit, and it is also the more difficult and

pertinent challenge. Thus, unless explicitly noted otherwise, references to Quine's challenge will be to the latter interpretation.

In reply to Wright (1981), which presents a version of Quine's challenge, Evans suggests that the challenge can be met by 'providing a causal, presumably neurophysiologically based, explanation of comprehension' (Evans 1981, 127). When such explanation is available, Evans claims, 'we can simply see' which theory is correct (*ibid.*). Evans goes on to suggest three additional plausible types of empirical evidence for one of a set of extensionally equivalent grammars as tacitly guiding a speaker.⁶ First, empirical evidence for the theory we actually use could come from the patterns in which we acquire dispositions, and second, from the patterns in which we lose dispositions, perhaps due to linguistic impairment. Thirdly, evidence can be culled from our (empirically testable) perceptions of linguistic structure in sentence perception (Evans 1981, 127-29; cf. Chomsky 1986, 252-87, Larson and Segal 1995, 56-62). (A clever thought experiment designed to show such evidence could be misleading is in Davies 1987, 451-53.)

2.2 Do All Processes Involve Tacit Knowledge?

We turn now to a well-worn argument against any attempt to explain linguistic competence with tacit knowledge. The argument has more critics than defenders, though Searle has employed versions of it (Searle 1983, 262-272; 1984, 28-31, 47-50). It goes something like this:

Suppose you posit a cognitive state called tacit knowledge to explain linguistic competence. If the general line of reasoning for positing this state is sound, why can't we invoke cognitive states to explain digestion? Just as competent speakers cannot explain how they know which strings are meaningful and which are not, so too proficient digesters cannot explain how they alter their stomachs to appropriately digest some food and reject indigestible food. In short, they 'interpret' their digestible input correctly and 'judge' the indigestible input as not part of their dietary corpus. But since the ability to digest is not *cognitive*, we should not posit a cognitive state to explain it. *Mutatis mutandis*, we should not posit tacit knowledge of a semantic theory to explain linguistic abilities.

Discussions of versions of this argument are in, *inter alii*, Nagel 1969, pp.172-174, Fodor 1975, p.74, fn.15, Chomsky 1986, pp.239, 241, and Wright 1986, pp.41-43.

A primary response is to defend differing general structures of the best theories of linguistic competence and digestion: unlike digestion, the best theory of linguistic competence entails that 'a representation of the rules they follow constitutes one of the causal determinants of their behavior' (Fodor 1975, 74; cf. Chomsky 1986, 244, 253-257). Employing linguistic capacities produces or requires certain belief-like states, such as whether 'Sta nevicando' means that it's snowing, or whether a string is a sentence of one's language. For linguistically competent organisms, their competence involves such beliefs. (This is an empirical defense, and so it would not follow that such beliefs are constitutive of one's competence, only evidence for it.) On the other hand, there is no reason to impute beliefs to digestively proficient organisms as such. We can account for the ability to digest good food and reject bad food without positing beliefs, explicit or implicit. (Cf. Lepore 1996 for a discussion of the epistemological import of linguistic beliefs.)

Nagel offers additional support for tacit linguistic knowledge, which invokes consciousness. He argues that 'In the case of language-learning...conscious

apprehension of the data...is essential; and what the individual can do as a result of his linguistic capacity is to speak and understand sentences' (Nagel 1969, 174). He compares statements of a tacitly known grammatical theory to statements that express cognitive attitudes revealed by psychoanalytic techniques, and he suggests that what they share is that it is often possible (at least in principle) to evoke a sense of recognition in the subject of the correctness of the attribution of the belief (or other attitude), and that this recognition will be, as it were, 'from the inside' (p.176). Nagel's aim is to drive a wedge between phenomena like digestion and linguistic competence by urging crucial connections with consciousness for the latter which the former lack (although cf. Chomsky 1986, 230).

However, Nagel never clarifies why we should suppose consciousness plays a role in language acquisition or competence. The literature on formal learning theory contains numerous descriptions of algorithms that can 'learn' small fragments of natural languages. When proposed algorithms fail to converge on the correct language, the problem is not that the system implementing the algorithm lacks consciousness. To take an example, Gibson and Wexler's Trigger Learning Algorithm learns any grammar in a hypothesis space of languages defined by a few parameters, and does so simply by ('unconsciously') reacting to its own failure or success at parsing the current input string (Gibson & Wexler 1994).⁷ In this literature, 'learning' is a technical term, though the aim is to model human learning. Nagel can always reply that consciousness is crucial to the actual learning of a grammar by a human. However, since there are attempts to uncover what is needed for learning not requiring consciousness, further defense of the connection between consciousness and learning is needed before any connection can differentiate linguistic competence from digestion.

2.3 Do Speakers Really Know a Grammar?

In addition to asking what justifies positing a distinctively cognitive capacity to account for linguistic competence, one might wonder whether the capacity is knowledge. As noted earlier, there is a difference between typical cases of knowing and so-called knowledge of a grammar. Various philosophers argue for psychological differences between typical beliefs and the information bearing states constitutive of 'knowledge' of grammar, and that these differences rule out the latter as beliefs (so, *a fortiori*, as knowledge as well) (Evans 1981, 131-32; Wright 1986a, 33-34, 41-43; Stich 1978). We will focus on Stich (1978).

One difference is that typical beliefs are accessible to consciousness: Attention 'suitably directed to the content of the belief' leads to 'a certain sort of conscious experience' (Stich 1978, 504). You may not be thinking about how you brush your teeth, but, if asked, you will have a conscious episode that involves reflecting (perhaps in detail) about how you do so.⁸ On the other hand, if asked to articulate the semantics of 'every', or just the part that explains why 'Every plane landed together' is ill-formed while 'All the planes landed together' is not, you might not know. In fact, even if told why 'every' behaves this way, you still might not believe it (in some sense, at least).⁹

Secondly, a typical belief 'inferentially integrates' with other beliefs, but states carrying grammatical information need not. For typical beliefs, if a subject believes that *p*, and comes to believe that if *p*, then *q*, she will also come to believe that *q*. Similarly, for other common deductive and inductive inferential schemata.¹⁰ To some extent, 'beliefs' about grammar share this property. For instance, the state of 'believing' that predicative noun phrases obey rule *R* may be inferentially connected to one's explicit belief that 'He is stupid and a liar' is fine, but 'He is a liar and John' is not.¹¹ However, though grammatical 'beliefs' may enter into inferential relations, Stich's point is that it is nonetheless severely restricted as to what kinds of inferences they can enter into. So, suppose you 'believe' predicative noun phrases obey rule *R*, and you also explicitly believe (perhaps because a wealthy theorist told you so) that if predicative noun phrases obey rule *R*, you will receive a million dollars. Despite the ingredients for a simple *modus ponens*, you do not come to believe you will receive a million dollars. You don't, Stich suggests, because grammatical 'beliefs' do not enter into inferential relations with other beliefs in the 'promiscuous' ways typical beliefs can. Similarly, most of us never feel an incompatibility between a tacit belief and an obviously contradictory conscious belief (Stich 1971, 489).

The foregoing argument challenges whether linguistic information bearing states are beliefs, and also whether it *matters* if they are. To see this, note that Stich assumes that the relevant states represent a theory of the language, and they are causally efficacious in linguistic comprehension. Whether such states are 'subdoxastic' or full-fledged beliefs depends largely on how beliefs function. So what is achieved by endorsing a theory that requires that *Xs* are beliefs (cf. Stich 1978, 514-515)? Are we seeking the true nature of reality or of our concepts? Are we trying to develop a useful concept for cognitive science? Whether these states are beliefs might be important to someone like Dummett, who believes that a theory of meaning must explain how language use is rational (cf. the opening pages of Dummett 1975, Dummett 1976; 1978, 104; cf. also Smith 1992, 124-31, Wright 1986b, 215-216, and Lepore 1996, 50). If linguistic competence is located primarily in subdoxastic states, perhaps we should concede that it is 'outrageous' to suppose that the type of propositional attitude speakers bear to their grammar is knowledge, in the usual sense (McGinn 1981, 290). We might instead follow Chomsky invoking the term of art 'cognize' for a sort of propositional attitude speakers bear to grammars (Chomsky 1986, 265-69).¹² (As a point of procedure, we will use the traditional 'tacit or 'implicit' knowledge, with no presumptions as to the nature of the type cognitive state it is. If you doubt such states are knowledge, treat our uses as privative adjectives, as McGinn suggests (McGinn 1981, 290).) A principal way to justify that speakers cognize grammars continues to be that assuming so better explains linguistic competence than any other hypothesis.

2.4 *Dispositionalism: Two Alternatives to Cognitivism.*

In this section, we will sketch two alternatives to cognitivism, what we shall call **unstructured** and **structured dispositionalism** (UD and SD, for short). We shall begin with UD. The cognitivist supposes that the hypothesis of tacit knowledge of a grammar is part of the best theory of linguistic competence, and so she posits tacit knowledge, thereby freeing herself to exploit any advantages of the hypothesis (as well as incurring its disadvantages). UD differs from cognitivism because it makes no strong claim about the relation between a grammar and a speaker. According to UD, a speaker may not tacitly know (or cognize) a grammar of her language. Its task is to construct a grammar that 'fits' (in Quine's sense) a speaker's dispositions to verbal behavior (where 'behavior' need not be understood in Quine's sense) (cf. Quine 1975).

UD is a *methodological* alternative to cognitivism, differing from it only about the scope of the project of devising a semantic theory for a natural language. Cognitivism requires a theoretical description of the semantic features of the target language that expresses the content of a representational state of the speaker which explains semantic competence. UD, on the other hand, requires a true semantic theory, but posits no psychological mechanisms. (At the other end of the spectrum is what we shall call non-cognitivism, according to which we lack tacit knowledge of a semantic theory.¹³ (We will return to this position below.)

Although UD is less bold than any account that purports to specify the psychological mechanisms that underwrite linguistic competence, its modesty also buys stability: a UD theory can be correct regardless of how a physical system like the human brain realizes dispositions constitutive of linguistic competence. Questions about realization are someone else's concern, perhaps the neuroscientist's. In this sense, then, the semanticist determines (in detail) the goal of what is an empirical problem for the neuroscientist and a design problem for the AI researcher. Furthermore, this naturally divides the theoretical work in accounting for linguistic competence. A UD defender might argue that cognitivism has semanticists strongly constraining the *architecture* of psychological and perhaps even neuroscientific theories. UD, on the other hand, only has semanticists constraining the *goals* of such theories. UD requires semanticists to inform psychologists about the semantic data to be explained, while cognitivism further requires semanticists to inform psychologists how to construct a theory that accounts for the data. Of course, the UD theorist is not suggesting that tacit knowledge posited by cognitivism is wrong; the essence of UD is quietism.

If UD is the correct methodological stance, why should finiteness concern us? One might object that the finite amount of our mental storage space, computational powers, and language acquisition time are all (strictly speaking) empirical hypotheses (cf. Davidson 1965). What justifies attention to these empirical data and not others? In response, first note that the dispositionalist is devising a theory to be used by the psychologist; he is not devising the theory used by a speaker. So, though the finiteness constraint is justified by the attention span of psychologists, it is also justified by the sorts of empirical data mentioned above. Although a UD theorist is

quiet about the nature of the psychogrammar (in George's sense), he needn't be completely oblivious – knowing basic finiteness facts about humans, he can try to respect this very modest empirical constraint. If other facts became as uncontroversial, they too might be incorporated into the dispositionalist's agenda. Perhaps, then, the rubric of dispositionalism houses a spectrum of theories, depending on how uncontroversial other data are.

In contrast to UD, which broadly characterizes dispositions to verbal behavior without a stance about which dispositional components comprise this larger collection (or how they do), **structured dispositionalism** (SD, for short) does take a stance. According to SD, corresponding to each axiom in a correct meaning theory is a unique disposition. Following Evans (1981), consider a finite language *L*, with ten proper names and ten one-place predicates, for a total of one hundred sentences. A speaker *S* has dispositions corresponding to a base clause (in a meaning theory for *L*) that says that 'a' refers to John just in case *S* has a disposition such that,

(2.5) For any quote-name Φ of any predicate of *L* and any predicate Ψ of the metalanguage of *L*, if *S* has the disposition corresponding to a clause that says something satisfies Φ iff it is Ψ , and *S* hears an utterance of the form $\Phi^{\wedge}a$, *S* will judge the utterance true iff John is Ψ .¹⁴

'connectedly', Evans writes, *S* has a disposition corresponding to the clause that says that something satisfies 'F' iff it is bald just in case *S* has a disposition such that,

(2.6) For any object *x* and any quote-name α of a name in *L*, if *S* has the disposition corresponding to the clause that says that α refers to *x*, and *S* hears an utterance of the form $F^{\wedge}\alpha$, *S* will judge the utterance true iff *x* is bald (Evans 1981, 124-25).

In addition to hypothesizing individuation conditions for dispositions that constitute a grammar, Evans recommends such talk to be understood in a 'full-blooded' sense: *S*'s dispositions are states of *S* appropriately causally responsible for the relevant patterns of behavior. Thus, SD posits a network of possibly non-cognitive dispositions constitutive of semantic competence. If they are non-cognitive (i.e., independent of any cognitive apparatus), SD is more than a methodological alternative to cognitivism. Despite using 'tacit knowledge', this is how Evans construes SD (cf. Evans 1981, 120-121, 124, and especially 133-134). On the other hand, SD may be a mere methodological alternative to cognitivism, if one is quietist about underlying the dispositional or categorical bases.

Wright raises three problems for SD. First, it is circular about understanding names and predicates: competence with a name is given in terms of competence with predicates, but competence with a predicate is given in terms of competence with names (Wright 1986a, 39-40; 1986b, 232-233). Secondly, when axioms are replaced with their corresponding dispositions, Quine's challenge remains: any empirical data that supports ascribing a set *D* of linguistic dispositions corresponding to a grammar can be made to support the ascription of a distinct set

D' of linguistic dispositions corresponding to an extensionally equivalent grammar, by exploiting 'appropriate hypotheses, of a *non-semantic* sort, about the presumed causal substructure' of the two sets of dispositions' (Wright 1986a, , 1986b, 231; cf. Davies 1987, 451-453). Finally, Wright notes that in the object language under consideration, it is natural to construct a compositional meaning theory using a compositional axiom, such as Evans',

(2.7) A sentence coupling a name with a predicate is true iff the object denoted by the name satisfies the predicate (Evans 1981, 123).

Wright then argues that a meaning theory would be 'crippled' without something like (2.7); but that SD need not postulate a disposition corresponding to (2.7). A speaker with the dispositions in (2.5) and (2.6) 'is thereby disposed to attach the proper significance to name-predicate coupling – since he is thereby disposed to attach the proper significance to sentences formed by coupling names and predicates' (Wright 1986a, 38; 1986b, 232). But now there is discordance between the details of the meaning theory and how SD says the meaning theory is 'realized'. If (2.7) is crucial to articulating a meaning theory, but its corresponding disposition is otiose in an account of linguistic competence, then the dispositions SD posits bear no simple one-one relationship to the axioms of the theory SD advertises. (For further discussion of dispositionalism and Wright's objections to SD see Davies 1987.)

2.5 *Semantic Non-cognitivism and Transductionist Theories.*

We began §2 with cognitivism and a battery of arguments against it. We turned to various forms of dispositionalism, which, to varying degrees, are alternatives to cognitivism. We turn now to another alternative to cognitivism, which we shall call **non-cognitivism**.

Strictly speaking, non-cognitivism is a form of dispositionalism, because it suggests that the best explanation of linguistic competence does not require cognitive relations to a semantic theory. As noted in §2.4, non-cognitivism entails that we lack tacit knowledge of a semantic theory. The standard way to support this entailment is to produce a theory which explains linguistic competence without appeal to tacit knowledge. Behind non-cognitivism is the idea that if linguistic competence can be so explained, then, assuming tacit knowledge does no theoretical work elsewhere, positing it is idle, and so, by Occam's razor, its existence should be denied. Although we will consider only one form of non-cognitivism, what we shall call **transductionism**, other types are available, such as those developed or suggested within a connectionist paradigm (cf. Elman, Bates, *et al* 1996, Rumelhart, McClelland *et al.* 1986, Langacker 1990).

To render knowledge of a semantic theory unnecessary, it suffices to show how competent users of a natural language could plausibly engage in the kinds of (linguistic and mental) activities they do without recourse to tacit semantic knowledge. Fodor articulates such a view in *The Language of Thought*, and still

endorses its relevant parts (Fodor 1975; 1990b; 1998; cf. also Schiffer 1987). We begin by sketching his position, and then turn to its criticism.

Transductionism. The main tenet of transductionism is that mental processing has the form of operations based on nomic properties of certain possibly complex mental objects. To be more precise, mental processing takes place because of operations on the syntactic features of expressions in a language of thought (LOT).¹⁵ On this view, the primary explanandum concerning natural language is how we communicate. According to transductionism, communication is the process whereby a sentence in a speaker's LOT, called a 'message' (Fodor 1975, 106), is mapped onto a phonetic string of English (say), which when produced in the vicinity of a hearer is in turn mapped onto (another token of) the message the speaker wished to communicate in the hearer's LOT. Other aspects of the hearer's processing algorithm (cf. §1.1) function to produce a belief about what the speaker said (i.e., a belief whose content is something like 'x said that P'). Successful communication lies in whether speaker and hearer share sufficiently similar transducing mechanisms between messages and heard strings (cf. p. 103). This is where transductionism becomes 'Gricean in spirit': expressions of a natural language like English acquire meaning in virtue of interpersonal similarities concerning the range of phonetic strings that can be used to communicate a given message (p. 104). However, this does not mean linguistics plays no role: a generative grammar for a natural language specifies for each message, 'the descriptions (morphological, phonological, syntactic, etc.) that a token [heard string] must satisfy if it is to conform to the linguistic conventions' for that natural language (p.109). Thus, one need not know (even though one surely does) that 'the dog' denotes the dog to be competent in English; one need only share with other speakers 'a knowledge of the descriptions that a written form must satisfy if it is to serve to communicate references to the dog to people who belong to that community' (p. 105).

According to Fodor, then, linguistic competence consists in an ability to map expressions of English onto correct expressions of one's LOT, and *vice-versa*, where correctness is a matter of conformity to the conventions of the community. Most interesting questions, such as 'What constitutes competence with respect to LOT?' and 'How do LOT expressions get their semantics?' are for the philosophy of mind and metaphysics (not epistemology and linguistics). Tokens of LOT get their meanings however they do, and have whatever meanings they have. Linguistic competence is just an ability to transduce objects of one sort (phonologically individuated strings) into objects of another sort (tokens of LOT). Fodor acknowledges this when he writes 'English *has no semantics*' (Fodor 1998, 9), other than whatever it inherits from the semantics of LOT. (A similar semantics-free view of linguistic competence is championed by, among others, Chomsky and Hornstein (Chomsky 1986, 1995, Hornstein 1984, 1988, 1989, 1991).

Fodor's view has been challenged by, among others, Lepore (1996). Lepore argues for epistemic consequences of linguistic competence that transductionism fails to explain. His point is that transductionism challenges the need to ascribe semantic knowledge by arguing that linguistic competence is constituted by a transduction relation between English and LOT. If someone hears you utter 'It's raining', she will reliably come to believe you said it's raining, because the transduction process from English to LOT is reliable, as well as are the other

'algorithmic' processes (cf. §1.1) needed to generate her belief. Lepore argues that if this were all there is to belief acquisition about what others say, transductionism would provide no account of one's own *reasons* for these beliefs. On the one hand, beliefs about what is said may be justified, at least on an externalist theory of justification, of the sort associated with, e.g., Goldman (1986). But on the other hand, transductionism provides no reason for why the interpreter acquires the particular belief she does about what you said. Compatible with transductionism, a speaker might be utterly 'clueless' as to how she acquired the belief that you said it's raining when you uttered to her 'it's raining', and she might also be clueless as to whether this belief is justified (Lepore 1996, 52). Following Davidson, Lepore suggests that 'nothing can count as a reason for holding a belief [about what's said] except another belief [about what the words uttered mean]' (p. 53; cf. Davidson 1986, 123), and that 'a belief that p (partly) rationalizes a belief that q only if the belief that p is (partly) causally responsible for the belief that q' (p. 53).

Although Lepore's argument is directed against transductionism, it also challenges various forms of dispositionalism. If the dispositions that constitute linguistic competence are non-cognitive, then although partly causally responsible for someone coming to have a belief about what another said, they cannot provide reason for one's having those beliefs.

III DOES KNOWING A GRAMMAR EXPLAIN LINGUISTIC COMPETENCE

3.0 *Is Modesty Enough?*

In this section, we will contrast modest and full-blooded meaning theories, and then review some objections to modest theories.

A **modest** meaning theory for a language L associates concepts with words and issues in meaning assignments to every sentence of L (cf. Dummett 1975, 102, 127; McDowell 1987, 72-73; Dummett 1987, 263-264; McDowell 1997, 119-120). Any theory that aims solely to derive theorems of forms (M) or (T) for every sentence S of L is modest,

- (M) S in L means that p
 (T) S is L is true iff p

where 'p' specifies the meaning of 'S'.

Dummett favors full-blooded theories over modest ones. The former not only associate words with concepts, but explain 'what it is to have the concepts expressible by means of that language' (Dummett 1975, 101). Where a modest theory might tell us only that something satisfies 'red' iff it is red, a full-blooded one 'explains...to someone who does not already have the concept' red what grasping the concept of red is. For more on modesty and full-bloodedness, see Dummett 1975, 102; McDowell 1987, 62; 1997, 105-106.

Why would anyone want more than modesty? Harman answers as follows,

[W]e might know that the sentence 'All mimsy were the borogroves' is true if and only if all mimsy were the borogroves. However, in knowing this we would not know the first thing about the meaning of the sentence,

(3) 'All mimsy were the borogroves' (Harman 1974, 6; our numbering; cf. also, Dummett 1975).

Knowing (3) is insufficient for understanding 'All mimsy were the borogroves' unless one already understands or has the concepts expressed by 'mimsy' and 'borogroves' (cf. Block 1986, 110, for a related argument). The theories under attack by Harman are modest theories, even though he couches his objection in terms of truth theories. Harman's objection (the **mimsy argument**, for short) is driven by an assumption that he and Dummett share, namely,

(D) A theory of meaning for a language L is a theory of understanding for L (Dummett 1975, 99).

Dummett emphasizes (D) (Dummett 1975, 99, 100-101; 1976, 69ff; cf. Smith 1992), 112. The role (D) plays in the mimsy argument is evident in its (schematic) reconstruction,

- (M1) A meaning theory for L must explain understanding sentences of L [from D].
- (M2) Modest theories do not explain understanding sentences of L.
- (M3) Nothing else about such theories (e.g., how they were constructed or justified) explains this understanding.
- (M4) ∴ [by 1,2, 3] Modest meaning theories are defective.

3.1 A Standard Reply.

Dummett and Harman both anticipate a reply to (D) and (M1)-(M4) that denies (M2) (Dummett 1975, 114; Harman 1974, 6). Modest theories explain understanding, because they are couched in a metalanguage the speaker understands (or at least she already has the concepts expressible in this metalanguage). So, a speaker's grammar will generate an interpretation of (3) only if her grammar has the axiom that something satisfies 'mimsy' iff it is mimsy. But a grammar with this axiom requires the speaker already to understand (or have the concept expressed by) the word 'mimsy'. Since a speaker's grammar can interpret (3) only if she already understands, or has the concept expressible by, 'mimsy', (M2) is false, and the mimsy argument is unsound.

Anticipating some such reply, both Harman and Dummett rebut that assuming prior understanding or conceptual grasp puts modest meaning theories on a par with translation manuals. A translation manual consists 'in the statement of an effective method for going from an arbitrary sentence of the alien tongue to a sentence of a familiar language' (Davidson 1973, 129). Translation theories *qua* of theories of

understanding have been criticized on the grounds that one can know a translation of every sentence of one language into another language without understanding any sentence in the former, and so, without understanding what any sentence means (Lewis 1970, 18-19; Davidson 1973). Imagine a manual in English that translates Greek into Latin. It will contain items like “*ἄνθρωπος*’ translates into ‘homo’”, “*ἵππος*’ translates into ‘equus’”, “*κλεπῶ*’ translates into ‘claudio’”. One could use this manual to interpret Greek only if one already understood Latin (and English). Similarly, urge Harman and Dummett, any modest theorist must be presupposing that a speaker already understands (or has the concepts expressible in) the language in which the theory is specified. Dummett and Harman rebut that were this presupposition correct, a translationist could make it as well. When explaining linguistic competence, presuming a translation manual can explain linguistic competence is incorrect.

3.2 Higginbotham's Reply.

Higginbotham 1989b, p. 165 contests (D) by arguing that a semantics for a language and a speaker's understanding of it can come apart. Consider Putnam's speaker who cannot distinguish elms from beeches. This speaker might fully understand his language, but his language might induce only a partial interpretation of ‘elm’ and ‘beech’. So, the speaker fully understands ‘beech’ and ‘elm’ in his idiolect, but what they mean is not what they mean in English, since in English their extensions differ (though cf. Burge 1979). Higginbotham suggests this is not how we think of reference.

Our words do refer to certain things...even when our knowledge of reference is incomplete. Moreover, it appears that incomplete understanding does not even prevent attribution of the same *concept* to the ignorant as to the learned. As we learn, we seem to come to know, or to know more fully, what things we refer to and through what concepts we refer to them (p.155).

Thus, he recommends we consider the language fully interpreted, with a speaker having only a partial grasp. If he is right, it is unclear whether a semantic theory ought to account for what one knows when one understands language, particularly if understanding a language despite is compatible with said deficiencies with respect to ‘elm’ and ‘beech’.¹⁷

So for Dummett a theory for L is correct only if its meaning theorems *explain* understanding L, whereas for Higginbotham it might be correct even without any such explanation (which he doubts it can (Higginbotham 1989b, 166)). Nonetheless, for Higginbotham knowledge of such a theory could constitute *partial* linguistic competence. Invoking partial constitution is supposed to support Higginbotham since it explains how speakers can use expressions they only partly understand: according to Higginbotham, one can know what ‘x carried out a leveraged buyout of y’ means and not know what leveraged buyouts are. No such explanation is available to Dummett, since he demands a meaning theorem to explain one's having knowledge of a homophonic meaning theorem. According to Higginbotham, then, there cannot be a fully explicit full-blooded theory of the sort Dummett envisages because our words have fixed meanings (your use of ‘beech’ doesn't have elms in its extension even if you cannot distinguish elms from beeches) despite our lacking the

appropriate understanding constitutive of a full-blooded meaning theory. To see why, it will be useful to discuss Dummett's attack on modest theories (1975, 105-108).

Dummett observes one can know 'the Earth moves' is true without knowing that the Earth moves. He calls the latter *knowledge of the proposition expressed*. With truth-conditional meaning theories, the goal is to explain knowledge of the proposition expressed by (3.1),

(3.1) 'The Earth moves' is true iff the Earth moves.

Knowledge of (3.1) is not disquotational, such as (3.2),

(3.2) 'The Earth moves' is true iff the Earth moves' is true.¹⁸

What must one know to know the proposition expressed by (3.1)? Dummett suggests one must know the meanings of its used component words. In a truth-conditional framework, this means knowledge of base axioms. Hence, by a similar line of argument, something besides knowledge of the truth of the axioms is required for knowledge of the propositions expressed by axioms. What could this something else be? If we suppose it to be knowing the truth of the axioms used in a derivation of (3.1), then we have started a regress. The additional knowledge must be of a different sort if it is to explain knowing the proposition expressed by (3.1). This final claim is the primary argumentative engine driving him to the conclusion that meaning theories must be full-blooded.

A crucial aspect of Dummett's position is that knowing the proposition expressed by a meaning theorem depends on knowing the propositions expressed by the axioms from which it is derived. So failure to understand a term like 'beech' amounts to failure to know which proposition is expressed by 'x satisfies 'beech' iff x is a beech'. However, *prima facie*, speakers have varying degrees of knowledge of the meanings of expressions; furthermore, over time, they may acquire increased degrees of knowledge of these meanings. So any account of partial knowledge that Dummett offers must account for these phenomena too. Either Dummett can argue against treating imperfect speakers as partially knowing a fully interpreted language or he can account for partial understanding. The latter must show how full understanding can be achieved, and be consistent with a theory of meaning being a theory of understanding. (If one only partially understands 'beech', what effect does this have on its meaning?) If he accepts partial knowledge of our language and adopts the latter approach, then the account still must be simpler than Higginbotham's, since Higginbotham provides a simple explanation of the phenomenon. Thus, stories involving complex structures of related propositions (e.g., structures that relate a proposition that amounts to total knowledge of a word to propositions that amount to partial knowledge – which themselves may have to be interrelated) will not work.

3.3 McDowell's Reply.

McDowell offers two arguments against Dummett (McDowell 1987). One supports modest theories directly, and the other indirectly by, in effect, assaulting the mimsy argument.

McDowell defends modest theories as such,

- (McD1) Meaning theories are modest or full-blooded.
- (McD2) They cannot be full-blooded.
- (McD3) ∴ They must be modest.

His second argument is only a bit more complicated,

- (McD4) An explanatory meaning theory (in Dummett's sense) must be full-blooded.
- (McD2) Meaning theories cannot be full-blooded.
- (McD5) ∴ No meaning theory is explanatory (in Dummett's sense).
- (McD6) There can be a correct meaning theory.
- (McD7) ∴ [denial of (M1)] Meaning theories need not explain our understanding (in Dummett's sense) of an object language.

Which feature of the mimsy argument McDowell's second argument challenges depends on what counts as explanation. If explanations are Dummettian, the second argument attacks either (D) (i.e., that a theory of meaning is a theory of understanding) or the inference from (D) to (M1), depending on how one understands 'understanding', an issue we discuss below.

Turning to (McD2), why reject full-blooded meaning theories? Dummett replies that though modest theories pair expressions with concepts, by failing to explain concept possession they fail to explain linguistic understanding (Dummett 1987, 258-60; McDowell 1997, 111-12). For McDowell this dilemma is false: the issue is not about explaining concept possession, but whether we can do so and still respect the constraint that a meaning theory be full-blooded. McDowell argues that one feature of this constraint concerns the sort of explanation of linguistic competence that is required by a full-blooded theory (McDowell 1987, 61). Full-bloodedness requires explaining what it is to possess concepts associated with words. So suppose we have a full-blooded theory for some language L. Understanding this theory must suffice for one previously unacquainted with L to come to understand L (Dummett 1975, 103-104; 1987, 265-266). But that a full-blooded theory must be in language, McDowell's argument runs, creates problems. First, by the response to the standard objection (in §3.1), the theory is on a par with translationist theories. Secondly, since the current move requires us to explain how one understands another language, progress on the task of explaining understanding a language is nil. Thus, McDowell seems to be using a version of the mimsy argument, one that attacks full-blooded theories and their demands on the explanatory work of such a theory. However,

endorsing any such argument does not prevent McDowell from attacking the other form of the mimsy argument, one which attacks modest theories. We will discuss this below.

McDowell responds that 'a proper theory of meaning for a language would be formulated "as from outside" content altogether' (McDowell 1987, 61). This requires that a full-blooded meaning theory not use expressions which specify or presuppose a specification of the contents (of words, expressions, utterances, thoughts, etc.).¹⁹ Although this is opaque, it appears that one has specified the content of an expression as from outside content altogether, if the specification does not include a use of an intensional context. This restriction prevents an explanation of possessing the concept 'square' along the lines: One has the concept *square* iff one is disposed to believe of all and only square things one encounters *that they are square*.²⁰

Thus, a full-blooded meaning theory must

- (i) explain, what it is to have concepts denoted by expressions in the language, and
- (ii) it must accomplish using a vocabulary that does not specify the contents of words, utterances, thoughts, etc.

But McDowell (and almost everyone else) also rejects behaviorism, so a theory that purports to explain concept possession in terms of 'outward behavior' is untenable (McDowell 1987, 65). This entails a further constraint on full-blooded meaning theories:

- (iii) the theory cannot be behavioristic (McDowell 1987, 63-65).

McDowell doubts any theory can satisfy (i)-(iii).²¹ One might try by ascribing tacit knowledge of a meaning theory (that avoids behaviorism) where such knowledge 'shows itself partly by manifestation of the practical ability, and partly by a willingness to acknowledge as correct a formulation of what is known when it is presented' (Dummett 1978, 96).²² However, McDowell notes that any such appeal guarantees that the meaning theory will be indeterminate: when all possible data are in, with every other relevant theory as precisely determined as can be, extensionally non-equivalent theories equally compatible with the data still exist (McDowell 1987, 66-67; 1997, 112-115). (Cf. George 1986 for the differences between underdetermination and indeterminacy.)²³ That is, no matter how much empirical data we have concerning e.g. the meaning of 'square', it can be accommodated equally well by theories according to which 'square' does not mean *square*. For suppose we hypothesize that the meaning of 'square' is *square*, because speakers of the object language are disposed to call only squares 'square'. This evidence is equally well explained by the hypothesis that speakers are disposed to call squares or pieces of mud from the bottom of the ocean 'squares'. Even if there *is* evidence that they are not so disposed, other Goodmanesque hypotheses compatible with the data will always be available (e.g., the disposition to use 'square' to pick out squares or numerals more than 1,000 digits long). (Cf. Goodman 1954, chapter 3.)

Regardless of how much evidence is available for positing tacit knowledge, it will be finite, and so infinitely many extensionally non-equivalent grammars that account for the data equally well will exist. (The same result holds even if there were (*per impossible*) an infinite amount of data.)²⁴ The *locus classicus* for problems of indeterminacy is Quine 1960.

To sum up: McDowell's arguments for modesty rely on (McD2). A full-blooded theory explains linguistic competence only if one can learn it without already understanding a language. This suggests that a full-blooded theory can be given 'as from outside' content' altogether, thus rendering full-blooded theories behavioristic. Finally, invoking tacit knowledge is no help, for to do so renders the theory unacceptably indeterminate. So, if meaning theories must be modest or full-blooded (a premise that aches to be clarified and challenged), they must be modest.

Furthermore, the mimsy argument fails, because if the explanatory task of a meaning theory is as Dummett says, then either no meaning theory is correct or the inference to (M1) is unsound. Since the former is implausible (though adopted in Schiffer 1987), the second must be adopted, which entails the unsoundness of the mimsy argument. On the other hand, perhaps one need not demand as much as Dummett about what suffices for explanation in (M1)-(M3). It may be that arguing to (M1) is legitimate, but one's alternative conception of an explanation is such that the justification for (M2) is thereby undermined. This seems to be McDowell's negative position regarding full-blooded theories. We will not discuss his positive view, but he does argue that the theorems of a modest meaning theory suffice to explain linguistic competence (and do so without incurring the indeterminacy of a theory that posits tacit knowledge of a full-blooded theory) (McDowell 1987, 67-70, 73-76; 1997, 116-119).

In conclusion, the diversity and difficulty of the replies we have reviewed show that mimsy argument to be anything but simple. It combines independently problematic issues including disquotational theories of truth, theories of truth as theories of meaning, lexical semantics, the structure and possession conditions of concepts, the nature of explanation, and the interface between one's psychogrammar and one's other capacities for the rational use of language. These issues are more fundamental than the mimsy argument because it can be understood only when these other issues are better understood.

IV SEMANTICS AND LINGUISTIC COMPETENCE

In this paper, we have discussed major issues concerning semantic competence. However, space prevents treating every relevant issue. We will conclude by merely mentioning three issues a longer paper on knowledge and semantic competence should discuss. (i) The concept of tacit knowledge was central in §2. A variety of analyses of this concept, and of the related concept of tacit belief, are in Lycan 1986, Dennett 1987, Kirsh 1990 and Crimmins 1992. (ii) The 'Kripkenstein' problem about whether past evidence can determine that we are currently following a rule (of grammar, for instance), and more specifically, whether there can be any fact of the matter about what we mean by our words. This problem first appeared in Kripke 1982, and a good overview of the problem can be found in Loar (1985). (iii) Quine's

'indeterminacy of translation' is often explained in terms of a speaker's ability to translate utterances from another language, though the translated language may be taken to be the translating language. In this latter situation, the problem purportedly shows that no single correct translation manual (or set of extensionally equivalent translation manuals) exists. The problem originates with Quine 1960; further discussion is in Root 1976 and Lepore 1977.²⁵

Kent Johnson and Ernie Lepore
Center for Cognitive Science
Rutgers University

NOTES

¹ For an argument that the distribution of negative polarity items cannot be characterized syntactically and must be characterized semantically see Ladusaw 1980 (cf. Higginbotham 1995a, 5-7).

² Cowper 1992, p. 2 reports Keyser as saying, "We are trying to figure out what it is that people *act as if* they know".

³ Quine suggests the relevant form of guidance is 'an intermediate condition, between mere fitting and full guidance in my flat-footed sense...' (Quine 1972, 442). Whether he's right is irrelevant here.

⁴ In the terminology of George 1989b, Quine's challenge is, 'What evidence selects one theory of a psychogrammar over another?'

⁵ In addition, Quine writes, '...the new doctrine of the grammarian's added burden raises the problem of evidence whereby to decide, or conjecture, which of two extensionally equivalent systems of rules has been implicitly guiding the native's verbal behavior' (pp. 443-44); 'The problem of evidence for a linguistic universal is insufficiently appreciated' (p. 446); 'The enigmatic doctrine under consideration says that one of these analyses is right, and the other wrong, by tacit consensus of native speakers. How do we find out which is right?' (p. 448).

⁶ It is not clear Evans intended to defend cognitivism. Nonetheless, his remarks may be so construed. The details of his position are taken up below.

⁷ See also Niyogi & Berwick 1996.

⁸ Stich's claim about typical beliefs' principled accessibility to consciousness is about what would (likely) happen were the subject and her situation normal. Unconscious beliefs of psychoanalytic theory do not count, because the antecedent is not satisfied, inasmuch as (we may suppose) some psychological mechanism interferes with ordinary processes leading from a belief to conscious awareness of it (Stich 1978, 505).

⁹ It may be that nobody knows why 'every' and 'all' distribute as they do. The example is from Christine Brisson's dissertation, 'Some Wider Consequences of Narrow Scope' (Linguistics, Rutgers University, 1998).

¹⁰ Evans agrees that inferential integration is *constitutive* of belief, 'To have a belief requires one to appreciate its location in a network of beliefs' (Evans (1981), p.132). He also ascribes it to Wittgenstein 1969, §141.

¹¹ Although Evans denies this point (1981, 133). It is hard to see how our linguistic competence could be explained by appeal to information bearing states that cannot interact with one another (assuming a relatively simplistic theory of individuation of the relevant

information bearing states) and, more importantly, could not produce further explicit or implicit beliefs. This would render tacit beliefs unable to explain, e.g., a speaker's coming to believe explicitly what a particular utterance means.

¹² We will remain silent about the relation between knowing and cognizing; Chomsky himself vacillates on the extent to which cognizing and knowing overlap; however, he is consistent about the unimportance of overlap for explaining linguistic competence (Chomsky (1980), Chomsky (1986), pp. 265-69).

¹³ It would be an interesting project to compare the notion of dispositionalism (and perhaps even some versions of non-cognitivism) with the notion of "knowledge how". Doing this would require developing a clear account of the cognitive structure of the latter notion, which would take us too far afield from the present project.

¹⁴ '∧' means 'concatenated with'.

¹⁵ A footnote of Fodor's on the syntax of LOT is relevant here, 'Any nomic property of symbol tokens...any property in virtue of the possession of which they satisfy causal laws...would, in principle, do just as well. (So, for example, syntactic structure could be realized by relations among electromagnetic states rather than relations among shapes; as, indeed, it is in real computers.)' (Fodor (1987), p. 156, fn. 5)

¹⁶ Lepore and Loewer (1981) respond to Harman by arguing that one no more needs to understand the metalanguage in which (A) is written to know what (A) expresses than Galileo needed to know English for him to have believed that the earth moves.

(A) 'La terra si muove' is true in Italian iff the earth moves.

They do not disagree with Dummett, however, that knowing (A) requires concepts of the earth and movement. Whether this excludes modest theories as theories of linguistic competence is a topic for the rest of this section.

¹⁷ There may be a way to reconcile Dummett and Higginbotham, because there are several ways to understand crucial terms both in (D) and in the argument in which (D) is employed (cf., Smith 1992 for extensive discussion of Davidsonian, Dummettian, and Chomskian interpretations of (D)). One might suppose Dummett has something special in mind by a theory of understanding: 'once we can say what it is for someone to know a language, in the sense of knowing the meanings of all expressions of the language, then we have essentially solved every problem that can arise concerning meaning' (Dummett 1975, 133).

¹⁸ This point is not unique to Dummett. Cf., Chomsky 1986, 266, Fodor 1968, 633-34.

¹⁹ The formulation in the text preempts appeal to contents in explaining concept possession. However, a weaker restriction is available: for any name ϕ of any expression of the object language, one cannot use ϕ in a content clause (i.e., in an intensional context) in an account of the possession of the concept denoted by ϕ . This permits using other kinds of content clauses in accounting for possessing the concept denoted by ϕ . Further restrictions on this second proposal are needed; how does Dummett avoid psychologism if possessing the concept denoted by ϕ is explained *via* a content clause containing a use of ψ , the possession of which is explained *via* a content clause containing a use of ϕ ? One might restrict the expressions that can occur in the content clause(s) that explain the possession conditions of the concept denoted by ϕ to those expressions taken to denote innate concepts, or to those expressions that have the possession conditions of their concepts explained 'earlier' in some recursively described hierarchy. It is not clear Dummett would take the first option, given his reluctance to develop his theory so that it becomes more than minimally answerable to empirical psychological hypotheses. However, if one had an acceptable means for defending some class of expressions as usable in content clauses in explaining concept possession, it might help with difficult cases, such as explaining theoretical concepts.

²⁰ The restriction does not prohibit using 'square' in accounting for possessing the concept *square*; cf. McDowell 1987, 62.

²¹ The core of McDowell's reply is that meaning theories cannot have properties (i)-(iii). But his attack is primarily about the compatibility of (ii) and (iii); he argues that any theory formulated 'as from outside' content must be behavioristic. If behaviorism is unacceptable, theories cannot be specified 'as from outside' content. But this undermines full-blooded theories only if they must be formulated 'as from outside' content, which has yet to be established. Thus, endorsing McDowell's argument does not require rejecting full-blooded meaning theories.

²² Two points are relevant here. First, the second part of Dummett's claim about how tacit knowledge might be partly manifested is false: if we have tacit knowledge of a meaning theory, there are many principles of this theory we are unlikely to acknowledge as correct when presented with their correct formulation. (For further discussion and examples, see §2.0.) Second, appeal to tacit knowledge of the present sort places an additional constraint on the formulation of meaning theories: where C is any concept expressed by an expression of the object language, explaining what it is to possess C must not use an expression that expresses C. Since the current suggestion is that one has a kind of knowledge of the theory, violating this restriction amounts to explaining possession of C by appeal to an epistemic state one has only if one already has C (McDowell 1987, 66).

²³ Is a theory's vulnerability to indeterminacy much of a criticism, since every theory suffers as such? McDowell believes his view is immune from indeterminacy because content is 'present in the words... [T]he thought (say) that some table-tops are square can be heard or seen in the words 'Some table-tops are square', by people who would be able to put their own minds into those words if they had occasion to do so' (McDowell 1987, 69).

²⁴ George glosses indeterminacy as follows: 'Where there is slack between observation and theory we have underdetermination, but slippage between total theory (all facts, known or unknown) and theory is indeterminacy. If any choice among the many present or future, explanatorily adequate, underdetermined theories of the world would leave unsettled the truth or falsity of linguistics' claims, then we cannot make sense of there being objectively correct evaluations of these' (George 1986, 489).

²⁵ Special thanks to Matti Sintonen and Barry Smith for their generous comments on earlier drafts of this paper. Ned Block (1986), 'Advertisement for a Semantics for Psychology', in Stich and Warfield (1994).

REFERENCES

- Burge, T.: 1979, 'Individualism and the Mental', in P. French, T. Uehling, and H. Wettstein (eds.) *Midwest Studies in Philosophy*, vol. IV, University of Minnesota Press, Minneapolis, pp. 73-121.
- Brandon, R. N. & N. Hornstein: 1986, 'From Icons to Symbols: Some Speculations on the Origins of Language', *Biology and Philosophy* 1, 169-189.
- Cappelen, H. and E. Lepore: 1997, 'On an Alleged Connection Between Indirect Speech and the Theory of Meaning', *Mind and Language*, 12, 278-296.
- Chomsky, N.: 1959, 'A Review of B. F. Skinner's *Verbal Behavior, Language*', in Block (ed.), *Readings in Philosophy of Psychology*, Vol. 1, HUP, Cambridge, 1980, pp.48-63.
- Chomsky, N.: 1965, *Aspects of the Theory of Syntax*, MIT Press, Cambridge, MA.
- Chomsky, N.: 1986, *Knowledge of Language*, Praeger, Westport, Conn.
- Chomsky, N.: 1995, 'Language and Nature', *Mind*, 104, 1-61.

- Cowper, E. A.: 1992, *A Concise Introduction to Syntactic Theory*, University of Chicago Press, Chicago.
- Crimmins, M.: 1992, 'Tacitness and Virtual Beliefs', *Mind and Language* 7, 240-263.
- Culicover, P. W.: 1997, *Principles and Parameters*, OUP, Oxford.
- Davidson, D.: 1965, 'Theories of Meaning and Learnable Languages', in Davidson, 1984, pp. 3-15.
- Davidson, D.: 1970, 'Semantics for Natural Languages', in Davidson, 1984, pp. 55-64.
- Davidson, D.: 1984, *Inquiries into Truth and Interpretation*, Clarendon, Oxford.
- Davidson, D.: 1986, 'A Coherence Theory of Truth and Interpretation', in Lepore (ed.), *Truth and Interpretation*, Blackwell's, Oxford.
- Davies, M.: 1987, 'Tacit Knowledge and Semantic Theory: Can a Five per cent Difference Matter?', *Mind* 96, 441-62.
- Dennett, D.: 1978, *Brainstorms*, MIT Press, Cambridge, MA.
- Dennett, D.: 1987, *The Intentional Stance*, MIT Press, Cambridge, MA.
- Dummett, M.: 1975, 'What is a Theory of Meaning?', S. Guttenplan (ed.), *Mind and Language*, 1974, OUP, Oxford, reprinted in P. Ludlow (ed.), *Readings in the Philosophy of Language*, MIT Press Cambridge, MA, 1997, pp. 129-155.
- Dummett, M.: 1976, 'What is a Theory of Meaning? (II)', in G. Evans and J. McDowell (eds.), *Truth and Meaning: Essays in Semantics*, Clarendon, Oxford, pp. 67-137.
- Elman, J., E. Bates et al.: 1996, *Rethinking Innateness*, MIT Press, Cambridge, MA.
- Fodor, J.A.: 1968, 'The Appeal to Tacit Knowledge in Psychological Explanation', *The Journal of Philosophy* 65, 627-640.
- Fodor, J. A.: 1975, *The Language of Thought*, Harvard University Press, Cambridge, MA.
- Fodor, J. A.: 1987, *Psychosemantics*, MIT Press, Cambridge, MA.
- Fodor, J. A.: 1990a, *A Theory of Content and Other Essays*, MIT Press, Cambridge, MA.
- Fodor, J. A.: 1990b, 'Review of Stephen Schiffer's *Remnants of Meaning*', in Fodor, 1990a.
- Fodor, J.A.: 1990c, 'A Theory of Content, II: The Theory', in Fodor, 1990a.
- Fodor, J. A.: 1998, *Concepts*, Clarendon, Oxford.
- Fodor, J. A. and E. Lepore: 1998, 'The Emptiness of the Lexicon', *Linguistic Inquiry*, 29, 269-288.
- George, A.: 1986, 'Whence and Whither the Debate Between Quine and Chomsky?', *The Journal of Philosophy*, 489-99.
- George, A. (ed.): 1989a, *Reflections on Chomsky*, Blackwell, Oxford.
- George, A.: 1989b, 'How Not to Become Confused about Linguistics', in George, 1989a, pp. 90-110.
- Gibson, E. and K. Wexler (1994), 'Triggers', *Linguistic Inquiry*, 25, 407-54.
- Goldman, A.: 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Goodman, N.: 1954, *Fact, Fiction, and Forecast*, Harvard University Press, Cambridge, MA.
- Graves, C. J., J. Katz, Y. Nishiyama, S. Soames., R. Stecker, and P. Tovey (1973), 'Tacit Knowledge', *The Journal of Philosophy*, 70, 318-330.
- Harman, G.: 1974, 'Meaning and Semantics', in M. K. Munitz and P. K. Unger (eds.), *Semantics and Philosophy*, New York University Press, New York, 1-16.
- Higginbotham, J.: 1983, 'The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics', *The Journal of Philosophy*, 100-27.
- Higginbotham, J.: 1985, 'On Semantics', in Lepore (ed.), *New Directions in Semantics*, Academic Press, London, 1987, pp.1-54.
- Higginbotham, J.: 1986, 'Linguistic Theory and Davidson's Program in Semantics', in E. Lepore (ed.), *Truth and Interpretation*, Blackwell, Oxford's, pp. 29-48.
- Higginbotham, J.: 1989, 'Elucidations of Meaning', *Linguistics and Philosophy* 12, 465-517.
- Higginbotham, J.: 1989b, 'Knowledge of Reference', in George, 1989a, pp. 153-174.

- Higginbotham, J.: 1993, 'Grammatical Form and Logical Form', in J. E. Tomberlin (ed.), *Language and Logic*, Atascadero, CA: Ridgeview, pp. 173-196.
- Higginbotham, J.: 1994, 'Priorities in the Philosophy of Thought', *Proceedings of the Aristotelian Society*, Supp. Vol. 20, 85-106.
- Higginbotham, J.: 1995a, 'Sense and Syntax', *An Inaugural Lecture delivered before the University of Oxford*, Clarendon, Oxford.
- Higginbotham, J.: 1995b, 'Some Philosophy of Language', in L. Gleitman and M. Liberman (eds.), *Language: Invitation to Cognitive Science* (2nd ed.), MIT Press, Cambridge, MA., 1995.
- Hintikka, J.: 1980, 'Theories of Truth and Learnable Languages', in S. Kanger and S. Ohman (eds.), *Philosophy and Grammar*, D. Reidel, Dordrecht, Holland., pp. 37-57.
- Holtzman, S. & C. Leich (eds.): (1981), *Wittgenstein: to Follow a Rule*, RKP, London.
- Hornstein, N.: 1984, *Logic as Grammar*, MIT Press, Cambridge, MA..
- Hornstein, N.: 1988, 'The Heartbreak of Semantics', *Mind and Language*, 3, 9-27.
- Hornstein, N.: 1989, 'Meaning and the Mental: the Problem of Semantics after Chömsky', in George, 1989a, pp.23-40.
- Hornstein, N.: 1991, 'Grammar, Meaning, and Indeterminacy', in Kasher, 1991, pp. 104-121.
- Kasher, A.: 1991, *The Chomskyan Turn*, Blackwell, Oxford.
- Kirsh, D.: 1990, 'When Is Information Explicitly Represented?', in P. Hanson (ed.), *Information, Language, and Cognition*, University of British Columbia Press, Vancouver, pp. 340-365.
- Kripke, S. A.: 1982, *Wittgenstein on Rules and Private Language*, Harvard University Press, Cambridge, MA.
- Ladusaw, W.: 1980, 'On the Notion Affective in the Analysis of Negative-Polarity Items', *Journal of Linguistic Research*, 1.
- Langacker, R. W.: 1990, *Concept, Image and Symbol*, Moutin de Gruyter, Berlin.
- Larson, R. and G. Segal: 1995, *Knowledge of Meaning*, MIT Press, Cambridge, MA..
- Lepore, E.: 1977, 'Reply to Professor Root', *Philosophical Studies*, 32, 211-215.
- Lepore, E.: 1996, 'Conditions on Understanding Language', *Proceedings of the Aristotelian Society*, 41-60.
- Lepore, E. and B. Loewer: 1981, 'Translational Semantics', *Synthese*, 48, 121-133.
- Loar, B.: 1985, 'Critical Review of Saul Kripke, *Wittgenstein on Rules and Private Language*', *Nous* 19, 273-280.
- Lycan, W.: 1986, 'Tacit Belief', in R.J. Bogdan (ed.), *Belief*, Clarendon, Oxford, pp. 61-82.
- Niyogi, P. & R. C. Berwick (1996), 'A Language Learning Model for Finite Parameter Spaces', *Cognition* 61, 161-193.
- Peacocke, C.: 1994, 'Content, Computation, and Externalism', *Mind and Language*, 9, 303-335.
- Quine, W.V.: 1960, *Word and Object*, MIT Press, Cambridge, MA..
- Quine, W.V.: 1972, 'Methodological Reflections on Current Linguistic Theory', in Davidson & Harman (eds.), *Semantics of Natural Language*, D. Reidel, Dordrecht, pp. 442-54.
- Root, M.: 1976, 'Speaker Intuitions', *Philosophical Studies*, 29, 221-234.
- Rumelhart, D., J. McClelland, et al.: 1986, *Parallel Distributed Processing, vol. 1*, MIT Press, Cambridge, MA.
- Schiffer, S.: 1987, *Remnants of Meaning*, MIT Press, Cambridge, MA..
- Segal, G.: 1989, 'Seeing What Is Not There', *Philosophical Review*, **XCVIII**, 189-214.
- Segal, G.: 1991, 'Defence of a Reasonable Individualism', *Mind* **C**, 485-494.
- Segal, G.: 1994, 'Priorities in the Philosophy of Thought', *Proceedings of the Aristotelian Society*, Supp. Vol. 20.

- Smith, B.: 1992, 'Understanding Language', *Proceedings of the Aristotelian Society*, **XCII**, 109-141.
- Smith, B.: 1998, 'On Knowing One's Own Language', in C. Wright, B. Smith, and C. Macdonald (eds.), *Knowing Our Own Minds*, Clarendon Press, Oxford, pp.391-428.
- Stich, S.: 1971, 'What Every Speaker Knows', *Philosophical Review*, **80**.
- Stich, S.: 1978, 'Beliefs and Subdoxastic States', *Philosophy of Science*, **45**, 499-518.
- Stich, S. and T. Warfield (eds.): 1994, *Mental Representation: A Reader*, Blackwell, Oxford.
- Wittgenstein, L.: 1969, *On Certainty*, Blackwell, Oxford.
- Wright, C.: 1981, 'Rule-Following, Objectivity and the Theory of Meaning', in Holtzman & Leich, 1981, pp. 99-117.
- Wright, C.: 1986a, 'How Can the Theory of Meaning be a Philosophical Project?', *Mind and Language* **1**, 31-44.
- Wright, C.: 1986b, 'Theories of Meaning and Speaker's Knowledge', in Wright, 1993, pp.204-38.
- Wright, C.: 1993, *Realism, Meaning and Truth* (2nd ed.), Basil Blackwell, Oxford.