

Meaning and Formal Semantics in Generative Grammar

Stephen Schiffer

New York University

The history of semantics in generative linguistics raises questions. By the time Chomsky's *Aspects of Syntactic Theory* was published in 1965 generative grammars were understood to have a semantic component in addition to a syntactic and phonological component, and it was assumed that a speaker's knowledge of her language required her to have tacit knowledge of a generative grammar of it. At that time the question that defined semantics in linguistics was the form that the internally represented semantic theory should take, and that is what the defining question was taken to be in Katz and Fodor's seminal 1963 manifesto, "The Structure of a Linguistic Theory," the first serious effort to do semantics in generative linguistics. Then around 1970 linguistic semantics took a curious turn. Without rejecting the claim that speaking a language requires tacit knowledge of a semantic theory of it, linguists turned away from the project of characterizing the nature of that theory in order to pursue instead the Montague-inspired project of providing for the languages we speak the same kind of formal semantics that logicians devise for the artificial languages of formal systems of logic. 'Formal semantics' originally signified semantics for formal languages devised for the mathematical study of formal systems of logic, but the expression now has a meaning akin to 'analytical philosophy' and signifies the Montague-inspired approach to the semantical study of natural languages. At the same time, many theorists—including many formal semanticists—recognize that the theories semanticists construct under the formal semantics rubric can't plausibly be regarded as theories of the kind needed to explain a speaker's knowledge of her language. The obvious question this bifurcation raises concerns the relation between, on the one hand, the psychologically explanatory semantic theories still thought to be needed but no longer the object of study in linguistic semantics and, on the other hand, the theories formal semanticists are concerned to construct. That question, I shall argue, becomes urgent when we understand the way in

which a psychologically explanatory theory would supersede even the best formal-semantic theory.

I. *The Generative Grammar Hypothesis*

A speaker of a natural language has the ability to understand indefinitely many sentences of her language that she has never previously encountered; indeed, her ability to understand *any* sentence of her language does not depend on her having a prior acquaintance with it. What explains this remarkable ability? One familiar answer emanating from MIT would require what I shall call the *Generative Grammar Hypothesis* (GGH):

The ability of a speaker of a natural language *L* to understand sentences of *L* requires her to have tacit knowledge of a generative grammar of *L*, that being a finitely specifiable theory of *L* that generates one or more syntactic structures for each sentence of *L* and interprets those structures both phonologically and semantically.¹

Of course, while the sentence one understands may be novel, neither the words that compose it nor the syntactical rules by which it's composed are novel, and it's commonly thought to be a platitude that one knows what a sentence means by way of knowing what its words and syntactic structure mean. If, as GGH tells us, one tacitly knows a generative grammar of one's language, then that would throw considerable light on how the platitude comes to be true. No doubt GGH derives much of its appeal from its ability to explain the platitude.

We can't hope to assess GGH other than relative to some way of understanding the connected notions of *sentence understanding* and *semantic interpretation* as they occur in GGH. GGH's relation to the (apparent) platitude about meaning suggests that

¹ In saying that GGH "emanates from MIT" I don't intend to be representing any particular linguist's views, let alone some time slice of Noam Chomsky. Rather, I take GGH to encapsulate a view that has a certain *prima facie* appeal and is close enough to what many theorists accept to sustain the conclusions I shall reach in this essay. The notion of a generative grammar utilized in GGH is, however, the view Chomsky accepted in his (1965).

understanding a sentence is knowing what it means, and that, consequently, it must be the job of the semantic component of a generative grammar to issue, for each sentence of the language, in a theorem that tells us what the sentence means. That is indeed true, as far as it goes, but it's helpful only to the extent that we already know in what knowing a sentence's meaning consists; it doesn't relieve us of having to ask what way of understanding *sentence understanding* and *semantic interpretation* has the best chance of verifying GGH. Before turning to that and related questions, however, I should say a little something now about the parts of GGH that for present purposes I will assume not to be contentious.

Linguists like to speak of a language user's "internal grammar"; but a grammar is a theory, and theories are abstract things that can't be in anyone's head or anywhere else. Better to speak, as linguists also do, of a person's internally *represented* grammar. The notion of *tacit knowledge* of a grammar is intended to make talk of an internally represented grammar somewhat more precise. One has tacit knowledge of a proposition p when one is in a subpersonal, or subdoxastic,² state that represents p , p is true, and the information-processing role of the state, *qua* representation of p , depends on p 's being true. Generative grammar, as conceived by Chomsky and, I hazard, the vast majority of those linguists who work within the Chomskian paradigm, is a branch of cognitive psychology, and that conception is understood to carry with it a commitment to what is called the *Computational Theory of Mind* (CTM)—the theory, with which I have no quarrel, that human information processing is computation in a *lingua mentis*, a brain's neural language of thought, which I'll call *Mentalese*. The notion of "tacit knowledge" to which GGH appeals is commonly understood to be a notion whose proper explication presupposes CTM. I shall accept CTM as a working hypothesis of this essay and will therefore assume that we process information in a neural system of mental representation, the "language of the brain," and that for a subpersonal state to represent p it must contain a token of a Mentalese sentence that represents p . Consider, for example, the information-processing sequence that begins with a hearer's auditory perception of the utterance of a sentence and terminates in her knowing what the speaker meant in producing the utterance. The picture implied by GGH is that crucially implicated in that

² For 'subdoxastic', see Stich (1978), and for 'subpersonal', see Dennett (1969).

processing sequence will be two intermediate states: first, a state that represents the generative grammar of the language (or part of a language) the hearer shares with the speaker, and at some point after that a state that represents the sounds produced by the speaker as tokening a sentence with syntactic structure Σ and semantic interpretation M , for some structure Σ generated by the grammar's syntactical component and some interpretation M that the grammar's semantic component assigns to Σ . This is another intuitively attractive picture, for an apparent platitude correlative to the one already mentioned is that one knows what a speaker means in uttering a sentence at least partly on the basis of knowing what the sentence he uttered means, and how could one compute the meaning of a sentence unless one's processing accessed a grammar that generated its meaning?

Much of this essay will be concerned with the question of what must be true of a semantic theory in order for it to be the semantic component of a generative grammar that verifies GGH. Now, there are broader and narrower ways of talking about "generative grammar," and some of these broader or narrower ways may be fueled by questionable stances on substantive theoretical questions. To take one prominent example,³ Chomsky identifies a person's I-language with the internally represented generative grammar that results when the parameters in the *universal grammar* (UG) that is innately represented in the mind/brain of every normal child are caused by the very young child's exposure to the speech of those in her linguistic environment to take on the specific values that determine her to have out of all the possible languages she might have acquired the particular language she does acquire. But the information that can be used to fix the language-determining parameters of UG are, according to Chomsky, rather severely constrained by the nature of the *language faculty* (LF) and can't be information that can be acquired only by interactions between outputs of LF and the outputs of other cognitive and motor faculties. Consequently, the properties assigned to an expression by what he would call a generative grammar won't involve relations between the expression and external-world objects and properties, properties that would be needed to determine the expression's extension and enable speakers to use the expression in communication, but are instead "internalist" syntactic properties—albeit syntactic properties Chomsky

³ See e.g. Chomsky (1995).

pointedly calls “semantic” without denying that they are also syntactic—which explain how expressions with those properties are able to interact with other cognitive faculties in ways that enable speakers to have and utilize tacit knowledge about the properties their expressions acquire as a result of their use in communicative and other behavior, properties that are essential to what we pre-theoretically might regard as essential to language understanding. From the perspective I shall adopt, all that is simply verbal gerrymandering. Once it’s decided what we want to regard as “sentence understanding” and “semantic interpretation,” we can ask what sort of semantic theory, if any, the explanation of those things requires. It wouldn’t matter whether some of what must be tacitly known results from interactions between the language faculty and other cognitive faculties.⁴

Back, then, to the question of what sort of semantic theory provides the best bet for verifying GGH. As I remarked above, Jerry Katz and Jerry Fodor’s “The Structure of a Semantic Theory,”⁵ published in 1963, was the first attempt to do semantics in generative linguistics. ‘Semantic theory’ in the title of their article refers to the semantic component of what Katz and Fodor called a “synchronic description” of a language, where such a description “seeks to determine what a fluent speaker knows about the structure of his language that enables him to use and understand its sentences.”⁶ When they wrote their article, generative linguists, following the lead of Chomsky, used the expressions ‘generative grammar’ and ‘generative syntax’ more or less interchangeably, but as a result of Katz and Fodor’s article and the publication a year later of Katz and Paul Postal’s *An Integrated Theory of Linguistic Descriptions*,⁷ ‘generative grammar’ came to be used in the way I have used it in stating GGH. Thus, the aim of Katz and Fodor was to describe the semantic component of a generative grammar of a natural language that would be implicated in explaining a language user’s knowledge of that

⁴ Chomsky seems actually to agree with this, notwithstanding his well-known skepticism about the possibility of the sorts of truth-theoretic and communication-based theories philosophers would call semantic. See op.cit.: 24.

⁵ Katz and Fodor (1963).

⁶ Ibid., p. 171.

⁷ Katz and Postal (1964).

language, that is to say, in explaining how speakers of the language are able to understand its sentences.

It is to their credit that Katz and Fodor begin their article by recognizing that the question of in what understanding a sentence—or, equivalently, knowing its meaning—consists is a substantial question that does not admit of an off-the-cuff answer. One might, however, suppose that the question *does* admit of an off-the-cuff answer. After all, ‘meaning’ is a count noun, as apparently true sentences such as

- ‘Bank’ has two meanings
- The meaning of ‘cat’ in English = the meaning of ‘chat’ in French
- He doesn’t know what ‘pulchritude’ means

bear witness, so shouldn’t we suppose that if an expression *e* has meaning, then there is something *x* such that *e* means *x*, and, consequently, that to know what *e* means requires knowing that *e* means *x*? Katz and Fodor don’t consider this question, but if they had considered it they might have said something like the following:

If an expression *e*’s having meaning consists in there being an *x* such that *e* means *x*, then the true proposition *that e means x* would be there for one to know, and knowing it would be what is required for one to know what *e* means. Yet it seems clear that knowing what an expression means—at least in the sense of knowing what an expressions means in which we do know what expressions mean—doesn’t require knowing a proposition that specifies what it means. That is evidently demonstrated by the fact that, while every fluent speaker of English knows what the word ‘of’ and the sentence ‘Is she there yet?’ mean, it’s highly doubtful that anyone knows what they mean by virtue of knowing the propositions expressed by completions of, respectively,

‘Of’ means ...

and

‘Is she there yet means?’ means ...

While it isn’t clear what exactly we may conclude from this, it is clear that even if in some sense of ‘meaning’ we may equate understanding a sentence with knowing its meaning, we can’t assume we know in what knowing a sentence’s meaning consists.

Anyway, what Katz and Fodor did conclude was that understanding a sentence was the ability to determine “the number and content of the readings of a sentence, [to detect] semantic anomalies, [and to decide on] paraphrase relations between sentences....”⁸ And given that that is what the ability to understand sentences is, the semantic theory they claimed was needed to explain that ability, and thus to be the semantic component of a generative grammar that verifies GGH, must have two components: (i) a dictionary (“[it’s] certainly true that one component of a semantic theory of a natural language is a dictionary of that language”⁹) and (2) “projection rules” that “select the appropriate sense of each lexical item in a sentence in order to provide the correct readings for each distinct grammatical structure of that sentence.”¹⁰

It wasn’t long before heavy flak hit the theory Katz and Fodor went on to construct around those two components. The problem critics found with the theory wasn’t that the theory wasn’t *true*; it was rather that it wasn’t in any suitable sense a *semantic* theory. For a theory that offers dictionary definitions of a language’s words and then attempts to derive readings of its sentences on the basis of those definitions isn’t telling us anything about the *semantic properties* of the language’s expressions: it isn’t telling us what any expression means or refers to, or what the truth conditions of any of its sentences are. For in the doubtful event that any word even has a correct dictionary entry, that entry wouldn’t *per se* tell us what the defined word means; it would tell us only that it means the same as the dictionary’s entry for the word. A person can learn what a word means from a dictionary entry only if she already knows what the dictionary’s entry for the word means; a monolingual speaker of Arabic could know

⁸ Katz and Fodor, op. cit., p. 176.

⁹ Ibid., p. 181.

¹⁰ Ibid., p. 196.

every proposition asserted by a correct English-Italian dictionary and not know what any word in either English or Italian means. By Katz and Fodor's own lights, the semantic theory that functions as the semantic component of a generative grammar of a language must be able not only to give the "readings" for each sentence of the language, but also the *content* of those readings. So even if the theory they went on to construct were true, it wouldn't explain what they think needs to be explained.

The original critics of Katz and Fodor's theory focused on the inadequacy of the theory as a semantic theory, but one can also object to Katz and Fodor's construal of what it is to understand a sentence, and thus their take on what the semantic component of a generative grammar needs to explain. For there is one very obvious constraint that ought to be recognized by any theorist who seeks to say in what the ability to understand novel sentences of a language consists: whatever that ability is, it is one fully possessed by every fluent speaker of the language. Therefore, for no Φ can it be correct to say that the ability to understand novel sentences of a language is, or consists in, the ability to Φ unless the ability to Φ is possessed, and possessed to an equal degree, by every fluent speaker of the language. But the ability to determine "the number and content of the readings of a sentence, [to detect] semantic anomalies, [and to decide on] paraphrase relations between sentences" is an ability that some fluent speakers don't have at all and others have to considerably different degrees. At the same time, there is at least one ability that is clearly possessed by every fluent speaker of a language and that does seem to be entailed by the ability to understand novel sentences. This, suitably qualified, is the fluent speaker's ability to mean things by uttering novel sentences of her language and to know what others mean when they utter them. That is a thought to which I will presently return.

It didn't take long for it to become the consensus among linguists and philosophers of language that, as instructive and ingenious as it was, the first serious attempt to do semantics in generative linguistics was a step in the wrong direction. Linguists interested in semantics in the mid-sixties were casting about for a way to pursue their subject. They were primed for an external influence.

II. Formal Semantics

The external influence that set linguistic semantics on its current course was the model-theoretic approach to the metalogic of formal systems of logic that emerged from the work of logicians Löwenheim, Skolem, Gödel, Tarski, Church, Kripke and others in the years between 1915 and 1965, and the circumstances that account for how the model theory of formal systems came to influence the way linguists do semantics for natural languages occurred in Los Angeles in the late sixties. Let me explain.

By 1965 it was taken as established in the metalogic of formal systems of logic that mathematically rigorous precisifications of the pre-theoretic notions of *logical truth* and *logical consequence* could be given in terms of the explicitly defined notion of a formula of a formal language being true under a model-theoretic interpretation of the language. The value of such definitions was that they made it possible to give mathematically rigorous proofs of the soundness and completeness of systems of logic containing the formal languages for which the model-theoretic definitions could be given, when soundness and completeness were themselves defined in terms of the model-theoretic account of logical consequence.

Also prevalent in 1965 was the view that “logical methods of formal language analysis did not apply to natural languages.”¹¹ This had been Frege’s and Russell’s view of natural language; it was the view of Tarski, who showed how an unproblematic notion of truth could be defined for certain formal languages but not for natural languages, which languages he dismissed as “inconsistent;” and it was the view not only of formally-minded philosophers such as Carnap and Quine but also of the “ordinary language” philosophers who opposed them, such as the later Wittgenstein, Ryle, Austin, and Strawson. It was in the climate of this prevailing attitude that in the late sixties, Richard Montague, a logician in UCLA’s philosophy department who had been a student of Tarski’s and whose work up to then had been primarily in intensional logic, turned his attention to the syntax and semantics of natural languages and wrote a trio of papers¹²

¹¹ Partee (2011; p. 18). Partee wrote this in commenting on what she called the “Ordinary Language vs. Formal Language War.”

¹² Montague (1970a), (1970b), and (1973).

whose bold motivating idea was provocatively announced in the first two sentences of the first of those papers:

I reject the contention that an important theoretical difference exists between [the artificial] formal [languages of logicians] and natural languages I regard the construction of a theory of ... truth under an arbitrary interpretation ... as the basic goal of serious syntax and semantics¹³

Montague's work would very likely have gone unnoticed by linguists had it not been for the coincidence of Barbara Partee's being an assistant professor in UCLA's linguistics department when Montague was working out his views and when David Lewis, with whom Partee had been an undergraduate at Swarthmore, was an assistant professor in UCLA's philosophy department and Hans Kamp was writing his dissertation under Montague. Writing about those days, Partee recalled that "one day David Lewis told me that Montague ... was starting to apply his work in logic to the semantics of natural language."¹⁴ Partee immediately became interested and began attending Montague's seminars. She initially found Montague's work difficult to understand, but she had Lewis and Kamp to help her understand it, and before long she came to believe that Montague's formal approach to semantics could be combined with Chomsky's formal approach to syntax to show the form that a semantic theory for a natural language should take. Partee brought Montague's ideas to linguists at a time when generative linguistics was at a loss as to how to do semantics for natural languages, and she proved not only to be a lucid conduit for linguists to Montague's ideas but also an extremely effective proselytizer of those ideas. It was in large part due to her efforts as a writer and teacher that "formal semantics" became and continues to be the dominate approach to semantics within linguistics.

We need to separate Montague's conception of natural-language semantics from the conception of it he inspired in Partee and those whom she influenced. Montague makes three remarkable claims pertaining to natural-language semantics—namely:

¹³ Montague (1970a, p. 188).

¹⁴ Partee (2004), p. 00.

- (1) There is no important theoretical difference between natural languages and the uninterpreted formal languages of systems of logic.
- (2) It's possible to treat a natural language as an uninterpreted formal language and to construct for it a model-theoretic semantics of exactly the same metamathematical kind that a logician would provide for the formal language of a system of intensional logic that captured the logical entailments expressible in the natural language.
- (3) The construction of such a semantics should be the goal of any serious semantics for natural language.

These claims are incredible, in the literal sense of that term, and it's difficult to suppose Montague didn't have his tongue in his cheek when he made them.¹⁵

Re (1) [There is no important theoretical difference between natural languages and the uninterpreted formal languages of systems of logic.] Taken literally, this claim is stupefying. For suppose L is a formally specified uninterpreted language, say, a recursively defined infinite set of sequences of types of marks that have been labeled "the sentences of L ." Then Montague's first claim is tantamount to saying that there is no theoretically important difference between, on the one hand, what is required to combine L with a deductive apparatus to form a system of deductive logic that we hope to show is sound and complete and, on the other hand, what is required for L to function as a public language of interpersonal communication. The claim is stupefying because it's of course obvious that on any remotely relevant scale for weighing the importance of theoretical differences there are several very important theoretical differences between what is required for the two very different tasks. As regards equipping L to be the formal language of a system of logic, there is a pre-theoretic notion of logical consequence (entailment by virtue of the forms of the entailing and entailed sentences) that systems of logic aim to explain in a mathematically rigorous way. But that pre-theoretic notion—requiring, as it does, that if a sentence is a logical consequence of a set of sentences, then it's impossible for the former not to be true when all the members of the latter are—crucially involves our pre-theoretic notion of truth, i.e. the notion of truth expressed by

¹⁵ Or at least it would be difficult for me to suppose that had Barbara Partee not assured me that Montague never said anything with his tongue in his cheek.

ordinary language truth predicates ('true' for English, of course), and, as Tarski liked to emphasize, that notion "in conjunction with the normal laws of logic leads inevitably to confusions and contradictions."¹⁶ Thus, a mathematically rigorous and consistent theory of logical consequence needs a mathematically rigorous and consistent precisification of the pre-theoretic notion of logical consequence, and that is precisely what the technical notion of *truth in a model of a formal language* is meant to make available. But the semantic properties involved in the definition of that notion are the *only* semantic properties expressions of a formal language of logic are given, and all that they need to be given. The expressions of a formal language of a system of first-order classical logic, for instance, have neither meanings, denotations, truth-values nor truth conditions. They have no *absolute* semantic values at all but only the ones they have *relative* to the infinitely many models of the language to which they belong, and there they don't even get anything worth calling their *meanings* relative to those models. Of course, that is unproblematic, for those relativized extensions are the only semantic properties that need to be assigned to the language's non-logical constants in order for the language to do the *logical* work required of it; after all, it's not as though the language is intended for use in interpersonal communication.

The requirements for equipping *L* to be a public language of interpersonal communication are very different from the requirements for equipping it to be the language of a formal system of logic, and, since a huge part of the *raison d'être* of natural languages is as systems for enabling interpersonal communication, natural languages must meet those requirements. Now, to communicate in a language is for a *speaker* to mean something by uttering a sentence of that language. It's no accident that if I want to tell you that I'm thirsty I'm much more likely to utter the sentence 'I'm thirsty' than I am to utter the sentence 'The flamingoes are flying south early this year', and this because in any ordinary circumstances uttering the first sentence is a very reliable way for me to mean that I'm thirsty, while uttering the second sentence in any ordinary circumstances is no way at all. And just as obviously, what makes 'I'm thirsty' such an effective instrument for meaning that one is thirsty is the fact that it has the meaning it has in English, and whatever property of a sentence constitutes its having the meaning it has is,

¹⁶ Tarski (1956), p. 267.

it's the sort of property that must be ascribed to the sentences of L if those sentences are to function as instruments for performing acts of speaker-meaning. It is, to be sure, no easy thing to say what sort of things natural-language meaning properties are, or how they help to determine the other semantic properties of the expression tokens that have those meaning properties, but for many properties it's very easy to say that *they* can't be those meaning properties. For example, if we define model structures for our formal language L so that in one of those models, M , the sentence Fa is true just in case Fido is a dog, while in another one of those models, M' , Fa is true just in case Richard Nixon is dead, then clearly we haven't even remotely done anything to make it possible for a person to meaning anything by producing a token of that sentence.

Re (2) [It's possible to treat a natural language as an uninterpreted formal language and to construct for it a model-theoretic semantics of exactly the same metamathematical kind that a logician would provide for the formal language of a system of intensional logic that captured the logical entailments expressible in the natural language] It's doubtful that Montague really meant that it was possible to construct a coherent notion of model-theoretic truth for *all* of English, as opposed merely to certain *fragments* of English. That is so if for no other reason than that English contains its own truth predicate, and if Montague had it in mind to give a model-theoretic account of truth for an object language that contained its own truth predicate, he surely would have addressed his teacher Tarski's argument to show that that was impossible. Another problem with the claim is that the model-relative semantic values of variables and non-logical constants must all be members of sets, but a natural language is able to talk about things that aren't members of any set, and, as Hartry Field points out,¹⁷ this leaves open the possibility that what a model-theoretic semantics determines to be a logical truth (in the model-theoretic sense of 'logical truth') may not even be true. But let's suppose Montague's claim (2) is true. That would no doubt be of considerable interest for many reasons and would refute the view that logical methods of formal language analysis don't apply to natural languages. But so what? What would it tell us about meaning in a natural language? That question returns us to Montague's claim (3).

¹⁷ Field (2008), Chapter 2.

Re (3) [The construction of a theory of truth under an arbitrary interpretation should be the goal of any serious semantics for natural language] The conjunction of this claim with Montague's first claim, that there is no difference of any theoretical importance between natural languages and the formal languages logicians construct, may seem to imply that Montague thought that the relative semantic properties assigned in an account of model-theoretic truth for a natural language were the only semantic properties the language's expressions could have. But of course Montague knew that some English expressions had absolute semantic values—that, for example, the name 'Richard Montague' referred to him and the sentence '1 + 1 = 2' was true. In so far as he intended us to take what he said seriously, would he deny that that difference between formal and natural languages was of any theoretical importance? I think in a sense he would. He rarely even mentions absolute semantic properties, but in one place where he is discussing lexical ambiguity (a type of ambiguity which, he says, is "rather uninteresting, to be sure"¹⁸) he revealingly says:

The use of a [natural] language would ideally involve not only the determination of the collection of *all* models of the language (a determination sufficient for the *logical* notions, that is, logical truth, logical implication, logical equivalence), but also the specification of a particular, actual model; this would be involved in characterizing *absolute* truth (as opposed to truth *with respect to* a model).¹⁹

I believe Montague's point is this: Of course natural languages are unlike the uninterpreted formal languages of logical systems in having absolute semantic values. But that isn't a very interesting difference, because the accommodation of those absolute semantic values doesn't require a different kind of semantic theory than the kind of model-theoretic semantics devised for the formal languages of logical systems, since to accommodate those values for a natural language *L* once we have an account of truth with respect to a model of *L* merely requires specification of which particular model is

¹⁸ Montague (1970a), p. 209.

¹⁹ Ibid.

L's "intended" model, the model whose assignments determine the semantic values the expressions of *L* actually have. What, however, would Montague say about the *meanings* of natural language expressions? I don't know, but given that he was content to understand the objects of propositional attitudes to be functions from possible worlds into truth-values (and given that he seems not to have had any interest in the use of natural language in communication) I suspect his attitude would have been that in so far as meaning was anything worth theorizing about, it should be identified with the intensions determined for the expressions of intensional formal languages in their model-theoretic interpretations.²⁰

So much for Montague's conception of natural language semantics. What was the conception of it he inspired in Partee and other linguistic semanticists, and, more generally, what is the conception of it that survives today in what in linguistics is called "formal semantics"?

In "Formal Semantics: Origins, Issues, Early Impact," published in 2011, Barbara Partee writes:

Montague did not invent model-theoretic semantics; but it was through his work that the model-theoretic approach became more widely known and adopted among linguists, with far-reaching changes to the field of linguistic semantics. One of the central ideas, not novel with Montague, is that *truth-conditions and entailment relations are basic. These are minimal data that have to be accounted for to reach "observational adequacy" in semantics* [my emphasis]. That principle, inherited from the traditions of logic and model theory, is at the heart of Montague's semantics and is one of the defining principles of formal semantics.²¹

²⁰ On the one occasion that I met Montague—in Berkeley in the summer of 1968 or 1969—he told me that nothing of any importance had been done in the philosophy of language since Tarski's *Wahrheitsbegriff*.

²¹ Partee (2011), p. 28.

Natural language expressions have meaning, and semantics is supposed to be the branch of linguistics that studies that meaning. What I find most instructive about the passage just quoted is the part I've italicized, wherein the only semantic facts mentioned as being semantic facts that a semantic theory of a natural language in linguistics should take it upon itself to explain are facts about the *truth conditions* of and *entailment relations* between the sentences of that language. In fact, for present purposes it is enough if we understand “formal semantics” in linguistics to refer to any semantic theory of a natural language *L* whose primary aim is to use the possible-worlds model-theoretic techniques of intensional logic to provide a compositional theory of *L* whose theorems state the possible-worlds truth conditions of *L* sentences and, with the help of a deductive apparatus for *L*, the entailment relations—or at least the logical entailment relations—that obtain among those sentences. Many of the things I'm about to say about such semantic theories will also apply to other conceptions of natural language semantics, such as the conception presupposed by the claim that natural language semantics may be accounted for in a Davidsonian or any other kind of compositional *truth* theory for the language, or on such generalizations of the model-theoretic-truth approach as dynamic semantics.

There are various mutually compatible theories of a natural language that would count as semantic theories of that language. My announced concern is about the nature of a semantic theory that could function as the semantic component of a generative grammar capable of verifying what I'm calling *the Generative Grammar Hypothesis* (GGH). From that perspective, I have two questions concerning the sort of theory that would pass muster in the “formal semantics” tradition of linguistic semantics as the sort of compositional semantic theory of a particular natural language whose construction ought to be the aim of linguistic semantics. Suppose *T* is a true model-theoretic semantics for English whose theorems assign possible-worlds truth conditions to the truth-evaluable sentences of English, assign semantic interpretations to the non-truth-evaluable sentences of English in terms of those assigned to appropriately related truth-evaluable sentences, and, in conjunction with a deductive apparatus for English, entail all the logical-entailment relations that obtain among English sentences. I very much doubt that there could be a *true* theory of English that satisfied that characterization of *T*, but

suppose there could be and that *T* is such a theory. Then my two questions, relative to that supposition, are:

- (A) Could *T* be the semantic component of a generative grammar that verified GGH?
- (B) If not, what, if anything, does that tell us about the privileged position formal semantics now holds in linguistic semantics?

Ever since the seventies, when formal semantics began to take hold in linguistics, there have been discussions in the literature that bear fairly directly on (A). These have been discussions about the relevance of formal semantic theories to psychological theories of linguistic competence. This is hardly surprising, since semantics is a branch of linguistics and Montague's influence began to be felt in generative linguistics at a time when there was a commitment to GGH in the discipline's nearly universal acceptance of "the Chomskian view that linguistics as a whole is a branch of psychology."²² Yet Montague's semantics was thought at the time to be opposed to the Chomskian view. For example, Richmond Thomason wrote the following in his Introduction to the 1974 collection he edited of Montague's papers:

Many linguists may not realize at first glance how fundamentally Montague's approach differs from current linguistic conceptions.... According to Montague the syntax, semantics, and pragmatics of natural languages are branches of mathematics, not of psychology.... This view is a corollary of Montague's strategy of studying natural languages by means of the same techniques used in metamathematics to study formal languages. Metamathematics is a branch of mathematics, and generalizing it to comprehend natural languages does not render it any less a mathematical discipline.²³

²² Partee (1978), p. 1.

²³ Thomason (1974, p. 2).

This is wrong: there is no opposition between a semantic theory's being a kind of mathematical theory and its being part of a psychological theory. The syntactical theories Chomskians construct are equally "mathematical" and no more psychological theories than is the theory of general relativity. What makes the Chomskian's endeavor psychological is that her aim is to explain knowledge of a language L in terms of tacit knowledge of a generative grammar of L , i.e. of a certain theory of L being represented in the brains of those who know L . But the theory that is internally represented, a particular generative grammar of L , is not a psychological theory. The psychological theory is the theory that an internal representation of that non-psychological theory is part of what accounts for the L speaker's ability to produce and understand a potential infinity of L sentences. Nevertheless, we are still left with question (A), the question whether a true model-theoretic semantics of English could be the semantic component of a generative grammar that verifies GGH.

Some formal semanticists seem open to a positive response to (A). In 1978 Max Cresswell suggested that to know what a sentence means "is simply having the ability to distinguish between worlds in which it is true and worlds in which it is false,"²⁴ but in 1985 he acknowledged that what he wrote in 1978 was "no doubt simplistic, since in fact the problem of propositional attitudes shows that often we *don't* know truth conditions; for we may not know that ' $7 + 5 = 11$ ' has the same truth conditions as ' 56 is a prime'."²⁵ The first sentence of Irene Heim and Angelika Kratzer's classic textbook on formal semantics is their declaration that "to know the meaning of a sentence is to know its truth-conditions."²⁶ And in what was evidently a handout I found on the internet for the first session of an MIT spring 2003 course on "Semantics & Pragmatics," the instructor, whose name isn't given but whose approach to semantics is that of model-theoretic semantics, begins with the statement that the "goal of semantics is to properly characterize semantic competence." David Lewis, whose "General Semantics" was enormously important in the development of formal semantics, emphasized a distinction between "two topics: first the description of possible languages or grammars as abstract

²⁴ Cresswell (1978, p. 12).

²⁵ Cresswell (1985, p. 145).

²⁶ Heim and Kratzer (1998).

semantic systems whereby symbols are associated with aspects of the world; and second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or population. Only confusion comes of mixing these two topics.”²⁷ I have seen Lewis’s comment cited in support of the view that “abstract semantic systems” aren’t supposed to be part of an account of what it is to know a language, but that is to misunderstand Lewis. It’s true that Lewis was himself opposed to appealing to what was represented in “the brain’s language of synaptic interconnections and neural spikes”²⁸ to explain what it is for a language to be the language of a certain population, but it’s entirely consistent with the distinction he drew that “the description of the psychological and sociological facts whereby a particular ... abstract semantic system [i.e. a particular language] is the one used by a person or population” should require those whose language it is to have tacit knowledge of the generative grammar that defines that “abstract semantic system.” Other formal semanticists—for all I know, most of them—disown any Chomskian connection between formal semantic theories and theories of linguistic competence. Barbara Partee, for example, wrote:

I’ve made a move toward the position that a theory of truth and reference for a language is not a theory of understanding. It’s true that the claim [that we need to postulate an internally represented compositional semantics in order to explain language understanding] has a kind of slogan status for semantics much as talk of “creativity” does for syntax. But I would say the real argument for compositional truth-theoretic semantics is not that language users can *understand* indefinitely many novel utterances, but rather that there are indefinitely many *semantic facts* to be explained, indefinitely many pieces of basic data about truth-conditions and entailment relations among sentences

²⁷ Lewis (1983), p. 184.

²⁸ Lewis (1999a: 13).

of any given natural language.²⁹

Partee is right to disassociate her interest in model-theoretic truth theories from attempts to explain language understanding, for it should be clear that if there is a semantic theory that is capable of being the semantic component of a generative grammar for a natural language that verifies GGH, that theory isn't a model-theoretic truth theory for that language. Crucial to assessing GGH is how we understand what it is to "understand a sentence," and, as we noticed earlier, there is an obvious but very important constraint on how we can understand sentence understanding: whatever we take the ability to understand sentences of a language to consist in, it must be something all normal speakers of the language share. As Montague, Partee and others conceive of a model-theoretic semantics for a natural language, it would account for the logical entailment relations that obtain among the sentences of the language, but that is clearly not anything that needs to be accounted for by a theory that aims to explain the ability of those whose language it is to understand sentences of their language. To be sure, it's plausible that one who understands, say, the connective 'and' must know that if 'Roses are red and violets are blue' is true, then so is 'Roses are red' and 'Violets are blue', but that is hardly to say that a speaker of a natural language knows *all* the entailment relations that a correct model-theoretic semantics for the language would capture.

A theory that did account for the entailment relations among the sentences of a language would be doing something that was superfluous to explaining the ability to understand sentences of the language; but of course a theory with a component that was superfluous for explaining certain facts might nevertheless explain those facts. Might, then, the semantic component of a generative grammar that verifies GGH be a possible-worlds model-theoretic semantics for English whose theorems assigned to each truth-evaluable sentence of English the conditions tokens of that sentence must satisfy to be true in an arbitrary possible world? I think not. Trying to theorize about sentence meaning without considering the very intimate connection between sentence-meaning and speaker-meaning is like trying to understand the properties of a hammer without knowing what hammers are for. For whatever else it might include, the ability to understand indefinitely many novel sentences of English must surely entail the ability to

²⁹ Partee (1988), p. 49.

know what speakers of English mean when they utter those sentences. That doesn't mean that if someone utters a sentence you understand then you are guaranteed to know what he means. One must know a lot more than the meaning of a sentence to know what someone meant in uttering that sentence. For example, you hear someone say 'Where is he?' and you don't know to whom she is referring with 'he'; or you don't know whether she is asking *you* where the referred-to male is; or you don't know whether she was asking a rhetorical question; or But if you hear the utterance of a sentence you understand *and* know relevant other things, then you will be able to know what the speaker meant in uttering that sentence. Suppose that you and a mono-lingual speaker of Japanese hear a friend of yours, Jane, say to you 'Will you leave if Lester isn't here by noon?' It's very unlikely that the Japanese speaker will know what Jane meant in producing her utterance,³⁰ but very likely that you will, and the crucial difference between you and him will be that you know what the sentence means and he doesn't. Now, as we're about to see, if GGH is true and you know what a sentence σ means, then there is some property ϕ such that your knowing what σ means consists in your knowing, perhaps tacitly, that σ has ϕ . What sort of property might ϕ be such that knowing that σ has ϕ will enable you to combine that knowledge with other things you know to infer what the speaker meant in uttering σ ? "The necessary and sufficient conditions for an utterance of the sentence to be true in an arbitrary possible world" can't be the right answer. It certainly can't be the right answer for sentences whose utterances can't have truth-values, such as the sentence 'Is she a politician?', but let's be charitable about the formal semanticist's ability to account for the meanings of non-truth-evaluable sentences in terms of the possible-worlds truth conditions of truth-evaluable sentences, if those truth conditions give the meanings of truth-evaluable sentences. In any case, there are truth-evaluable sentences which show that knowing the meaning of a sentence can't be equated with knowing its possible-worlds truth conditions. Two such sentences are:

(1) I would be well adjusted if I'd had different parents.

³⁰ Very unlikely, but not impossible: sometimes the context of utterance makes it likely that someone who has no knowledge at all of the language to which the utterance belongs will know enough to be able to have a pretty good idea of what the speaker means, and thereby a pretty good, but not quite as good, idea of what the uttered sentence means, even if she has no idea of what language is being spoken.

(2) She married an ax-murderer.

For suppose Kripke is right that it's metaphysically impossible for a person to have derived from any other sperm and egg than the pair that produced him. Then tokens of (1) would evidently have the same possible-worlds truth conditions as any other counterfactual conditional whose antecedent is metaphysically impossible; so clearly, knowing what those conditions are would on its own put one in an extremely poor position to have the knowledge you have about what a speaker would mean in a literal and unembedded utterance of (1). The problem with (2) is interestingly different. In a possible-worlds model-theoretic semantics that has the best chance of assigning to sentences the possible-worlds truth conditions that tokens of those sentences would actually have, a token of (2) would be true in an arbitrary possible world w just in case the female to whom the speaker referred in his utterance of that token was married in w to someone who belonged to the set of ax-murderers in w . But for any arbitrary world w one could conceivably know of the set of ax-murderers in w whether or not any given individual in w belonged to that set even though one had no idea whether or not that person was an ax-murderer in w . This could be so if one's knowledge of the set of ax-murderers in w was by enumeration, so that while one didn't know whether or not Jones was married to an ax-murderer in w , one did know that in w Jones's spouse belonged to the set whose members were {Bob, Carol, Ted, Alice, Kit Fine, ...}, where that set was the set of ax-murderers in w ; or it could be so if one knew that in w Jones's spouse belonged to the set of people who had the gene encoding the neurotransmitter-metabolizing enzyme monoamine oxidase A, where that set was identical to the set of ax-murderers in w . Therefore, possible-worlds model-theoretic semantics for English can't be a theory that explains sentence understanding, and therefore can't be the semantic component of a generative grammar that verifies GGH. To be sure, this wouldn't have disturbed Montague, who aimed only to explain the logical entailments that obtain in a language. So if linguists who are formal semanticists aim only to tell us the form that a compositional semantics for a natural language must take if it's to make the right assignments of possible-worlds truth conditions to the language's sentences, why should they care if those compositional semantic theories can't be the semantic components of internally represented generative grammars? Well, once we see what the semantic

component of a generative grammar of a language would have to assign to the sentences of that language, we'll see more than one reason why they should care. We may even want to conclude that, even if the theories formal semanticists construct are true, formal semanticists aren't really doing what semantics in linguistics should require them to do. Let's turn then to what might reasonably be taken to be the correct account of what a semantic theory would have to be if it's to be the semantic component of a generative grammar that verifies the Generative Grammar Hypothesis.

III. *The Best Bet Theory*

Let's remind ourselves where we are. We are asking what a semantic theory must be like in order to be the semantic component of a generative grammar that verifies GGH. How we answer that question will depend on what we take understanding a sentence to consist in. Whatever constitutes the ability of a native speaker of a language to understand sentences of her language must be something shared by all native speakers of the language. That constraint on language understanding led us to see that Katz and Fodor were mistaken about the data their semantic theory needed to explain. A second constraint also emerged. The *raison d'être* of natural language, or at least a huge part of it, is as a device for interpersonal communication, and whatever else understanding a sentence requires, it requires the ability to understand utterances of it, the ability to know what a speaker who utters the sentence means in uttering it. Even when you say to me 'Withers once more the old blue flower of day' and I know that you are thereby suggesting that it's time for a cocktail, I would not have known what you meant in uttering the sentence if I hadn't known what the sentence means. It should be obvious that a person couldn't count as knowing a language if she couldn't communicate in the language or know what others were communicating when they used the language. Knowing a language requires the ability to mean things in uttering novel sentences of the language and the ability to hear the utterance of a novel sentence and know what the speaker meant in uttering it.³¹ Now, independently of GGH we should know that, if *L* is

³¹ Even if, like Chomsky, one thinks that the use of one's language in interpersonal communication isn't essential to its being one's language, we would still want to explain what must be true of it in order for it to function as a device for interpersonal

the natural language used by a population P and σ is a sentence of L , then for some property ϕ , σ 's having ϕ constitutes its meaning what it does in L and, by virtue of that, σ 's having ϕ is crucial to enabling members of P to know what a member of P means if she utters σ . What sort of property is ϕ and how does it figure into the explanation of the information processing that takes a speaker of L from her hearing an utterance of σ to her knowledge of what the speaker meant in uttering σ ? The truth of GGH would constrain what can count as a correct answer. For if GGH is true, then:

- a) ϕ is assigned to σ by the semantic component of the speaker of L 's internally represented generative grammar for L ;
- b) ϕ is compositionally determined by the semantic component as a function of the semantic properties it assigns to the morphemes from which σ is constructed;
- c) when a speaker of L hears an utterance of σ and knows what the speaker meant in producing that utterance, at some point in the processing that takes her from her perception of the utterance to her knowledge of what the speaker meant, the hearer (i) enters a state that constitutes her tacitly knowing that σ has ϕ , after which (ii) that tacit knowledge is conjoined with other tacit or non-tacit knowledge she has, and (iii) from that combined state of knowledge she tacitly infers (no doubt via an "inference to the best explanation") that the speaker meant such-and-such in uttering σ .

What must ϕ be like if GGH is true? Here is a thought experiment that suggests an answer.

communication; and if, as seems plausible, it's because one's sentences have the meanings they have for one that enables one to communicate in the language, then it would remain the case that whatever internally represented semantic theory was needed to explain sentence understanding would also have to explain the language's ability to function as a medium of interpersonal communication.

You see a note on which the sentence ‘Is she a politician?’ is written. You don’t know who wrote the note, to whom it was written, or the circumstances in which it was written. You were however told that the sentence was produced by an English speaker, S , who intended the note to be taken literally (as opposed, say, to metaphorically or sarcastically). What would you know about what S meant in producing the note? A plausible first thought is that you would know that the speaker was referring to a certain female and asking whether she was a politician. Similarly if, keeping all else the same, the sentence on the note had been ‘It’s snowing’, then you would know that S meant that it was snowing in a certain place to which she was implicitly referring. What this suggests is that to know the meaning of a sentence σ is for there to be a type of speech-act A and a form of content Ψ such that one knows that in a literal and unembedded utterance of σ the speaker is performing an act of kind A whose content is of form Ψ . For example, knowing the meaning of ‘Is she a politician?’ requires knowing that in a literal and unembedded utterance of it the speaker would perform an act of *asking-whether* whose content was of the form x_f is a politician at t_u , where ‘ x_f ’ holds a place for a specification of the female to whom the speaker referred with his utterance of ‘she’ and ‘ t_u ’ holds a place for a specification of the time of utterance. Similarly, to know the meaning of ‘It’s snowing’ is (to a first approximation) to know that in a literal and unembedded utterance of the sentence the speaker would be performing an act of *meaning-that* whose content was of the form *it’s snowing at l at t_u* , where ‘ l ’ holds a place for the location to which the speaker implicitly referred and ‘ t_u ’ is as before. It should be *obvious* that the intimate connection that obtains between sentence-meaning and speaker-meaning demands that, if there are such things as sentence meanings, sentence meanings *somehow* constrain what speakers can mean in uttering sentences with those meanings, and constrain what can be meant narrowly enough so that uttering the sentence is an effective way to communicate propositions that satisfy the constraint. If GGH is correct, then it must be that the A - Ψ pairs that do the constraining in the way indicated are assigned to sentences by the semantic component of the internally represented generative grammar, and thereby count as the meanings of the sentences to which they are assigned. (We may reasonably assume that Ψ is also what is typically assigned to σ when embedded in a larger sentence as σ ’s contribution to the form of

content assigned to the larger sentence; but there may be exceptions to what is typically the case, as there would be for Frege, who was committed to a sentence's having a different meaning when it occurred in the complement clause of a propositional-attitude sentence.) Now, of course I haven't *shown* that the theory just adumbrated is the best bet for being the semantic component of a generative grammar that verifies GGH, but I think that it is³² and, more importantly, that it comes close enough to being the best bet to show that any theory that would verify GGH if taken to be the semantic component of a generative grammar for a language would supersede a correct model-theoretic truth theory for the language in just the ways that I will presently argue the adumbrated theory would supersede it. In any case, I will assume as a working hypothesis that the sort of meaning theory just sketched is the best bet for being the semantic component of a generative grammar that verifies GGH, and I will refer to it as the *Best Bet theory* (or simply as *Best Bet*).

For the formal semanticist, propositional speech-act sentences are just one more class of sentences for which she must provide truth conditions. For the Best Bet theorist, however, the semantics of those sentences is of foundational importance. This is because she won't know how to understand the content forms her semantic theory must assign to sentences before she understands what the content-ascribing clauses of propositional speech-act reports contribute to the truth conditions of those reports. It is, however, no easy task to know what the italicized clauses in sentences such as the following are up to semantically:

Al said *that Betty was a politician*

Al asked *whether Betty was a politician*

Al asked Carl *where Betty went on her vacation*

Al told Carl *to hire Betty*

Rather than embarking on the book-length project of trying to set out and defend a semantic theory of these sentences, I propose to help myself to the following two shortcuts.

My first shortcut is that I will *pretend* that assertoric speech acts—*meaning that such-and-such, asserting that such-and-such, telling so-and-so that such-and-such, ...*—

³² See Schiffer (2003, Chapter 3).

are triadic relations between speakers, Russellian propositions and times, although to simplify the exposition I will ignore tense-induced temporal references. Russellian propositions are structured entities whose building blocks are the objects, properties and relations our speech acts may be about, and I will follow the usual practice of representing Russellian propositions by ordered pairs of the form $\langle \langle x_1, \dots, x_n \rangle, R^n \rangle$, where $\langle x_1, \dots, x_n \rangle$ is an n -ary sequence of things and R^n is an n -ary relation (properties are 1-ary relations), and where $\langle \langle x_1, \dots, x_n \rangle, R^n \rangle$ is true just in case $\langle x_1, \dots, x_n \rangle$ instantiates R^n , and false just in case $\langle x_1, \dots, x_n \rangle$ doesn't instantiate R^n . So, if Al means that Betty is a politician, then (following the usual practice of dropping pointed brackets for one-membered sequences and, as I warned, ignoring the time parameter) we may represent the Russellian proposition Al means as $\langle \text{Betty, the property of being a politician} \rangle$. In some cases—for example, propositions denoted by ‘that’-clauses containing quantifiers or modal terms—while there should be no obstacle in representing those propositions as Russellian propositions, difficult theoretical questions must be resolved before we know how exactly to represent them as Russellian propositions. It's not important for my purposes that those issues be resolved, for the points I need to make can be made with respect to those Russellian propositions that are “singular propositions,” such as the proposition that Betty is a politician ($\langle \text{Betty, the property of being a politician} \rangle$) or the proposition that Harold loves Betty ($\langle \langle \text{Harold, Betty} \rangle, \text{the love relation} \rangle$). My assuming for the purposes of this paper that assertoric contents are Russellian proposition counts as *pretense* because I believe that those contents are *not* Russellian, for I see no acceptable way of solving Frege problems on the assumption that the assertion one would make in a literal utterance of ‘Samuel Clemens was Mark Twain’ is identical to the assertion one would make in a literal utterance of ‘Samuel Clemens was Samuel Clemens’, or that the assertion one would make in a literal utterance of ‘Fiona is fond of groundhogs’ is identical to the assertion one would make in a literal utterance of ‘Fiona is fond of woodchucks’. Now I do believe that assertoric speech acts are relations to propositions, in the generic sense of ‘proposition’ that means *abstract mind- and language-independent entities that have truth condition which they have both essentially and absolutely (i.e. without relativization to anything else)*. But a statement of Best Bet on the assumption that the relevant propositions are Russellian makes for a simpler and

more accessible exposition than the one that would result if the theory were stated using more plausible candidates for being propositional contents, and, as I trust will become clear in due course, the critical claims I will make on the Russellian assumption will transpose to a version of the theory that uses any other notion of propositional content. That is why my pretense is a *shortcut*.

My second shortcut concerns what I *won't* say about the contents of non-assertoric speech acts. Taking assertoric speech acts to be relations to Russellian propositions commits me, I believe, to taking all propositional speech-act contents to be either Russellian propositions or things defined in terms of them, such as forms of Russellian propositions (e.g. propositions of the form $\langle x, \text{the property of being a dog} \rangle$). I believe that for the purposes of this essay I may assume without argument that that commitment is unproblematic on the pretense that assertoric speech acts are relations to Russellian propositions, and that the arguments I shall offer won't suffer from my mostly focusing on the meanings of those sentences that can be used, when speaking literally, to perform assertoric speech acts.

The thought experiment of a few paragraphs back suggests that to know the meaning of

(1) She is a politician

is to know that in a literal and unembedded utterance of (1) there is a female x (or so the speaker intends) such that in uttering (1) the speaker means that x is a politician. Given the assumption that assertoric speech acts are relations to Russellian propositions, this suggests that at a certain level of theorizing we may represent the meaning of (1) as the pair

(2) $\langle \text{t}, x_f \text{ is a politician} \rangle$,

where 't' stands for the act type *meaning-that* and ' $x_f \text{ is a politician}$ ' stands for the propositional form $x_f \text{ is a politician}$, where ' x_f ' holds the place for a female to whom the speaker refers in uttering the sentence, which is equivalent to saying that ' $x_f \text{ is a politician}$ ' stands for the property of being a proposition p such that for some female x to whom the speaker refers, $p = \langle x, \text{the property of being a politician} \rangle$. Given the Russellian pretense, if GGH is correct, then an internally represented generative grammar for English will issue in a theorem that pairs (1) with (2), thereby encapsulating the

information that in a literal and unembedded utterance of (1) there is a female x (or so the speaker intends) such that in uttering (1) the speaker means that x is a politician.

So, relative to ongoing assumptions and constraints, Best Bet deems (2) to be the meaning of (1), but not, however, because that can be discerned in the ordinary language use of ‘meaning’, but only in the sense that there is some relation R —arguably having to do with facts pertaining to the use of language in communication—such that for some type of speech act A and propositional form Ψ , two unambiguous sentences will be equivalent in meaning just in case they both bear R to $\langle A, \Psi \rangle$.

Meaning is a property of expression *types*; a sentence type’s meaning is a *constraint* on what a speaker can mean in uttering tokens of the type, and the meaning of any expression constrains the contribution the expression makes to the meanings of the expressions in which it occurs. Since theorists sometimes speak of e.g. the “metaphorical content” of a sentence token, it’s useful to have a notion of *semantic content* as a property of a sentence token that is defined in terms of the meaning of the sentence type of which the token is a token. When GGH is assumed, we may say that:

- (SC) A proposition p is a *semantic content* of a token σ_τ of sentence σ if for some propositional form Ψ ,
- 1) σ means $\langle \vdash, \Psi \rangle$;
 - 2) p is of form Ψ ; and
 - 3) the speaker meant p in uttering σ_τ .³³

(Note that [SC] stipulates only a sufficient condition for being a semantic content, although, it may be noticed, 1)-3) are separately necessary as well as jointly sufficient for being a semantic content of a token of a *truth-evaluable* sentence. [SC] offers only a sufficient condition because we should also want to speak of the semantic content of interrogative, imperative, and whatever other non-truth-evaluable sentences there may be. The purposes of this essay may be served without venturing to define the notion of semantic content for those cases.) [SC] may then be used to define (at least to a first approximation) the semantic contents of expression tokens that are components of unembedded truth-evaluable sentence tokens:

³³ Given [SC], we may define the Gricean’s favored technical sense of ‘say’ thus: S said p in uttering σ_τ iff p is the (or a) semantic content of σ_τ .

c is the (or a) *semantic content* of a token ε_τ of expression type ε if for some unembedded truth-evaluable sentence token σ_τ ,

1) $\varepsilon_\tau = \sigma_\tau$ (& $\varepsilon = \sigma$) & c = the semantic content of σ_τ ;

or

2) σ_τ contains ε_τ ;

3) c is ε_τ 's contribution to the (or a) semantic content of σ_τ .

Then let's say that an expression type has a *constant meaning* just in case all its tokens that have semantic contents must have the same semantic contents, and that an expression type has a *non-constant meaning* just in case different tokens of the type may have different semantic contents. So, if, as may seem plausible on the pretense that propositional contents are Russellian propositions, it's correct to say 'Hillary Clinton' means Hillary Clinton, 'politician' means the property of being a politician, and 'Hillary Clinton is a politician' means $\langle \uparrow, \langle \text{Hillary Clinton, the property of being a politician} \rangle \rangle$, then each of those expressions has a constant meaning, as Hillary Clinton will be the semantic content of every token of 'Hillary Clinton' that has a semantic content (*qua* name of the former US Secretary of State), the property of being a politician will be the semantic content of every token of 'politician' that has a semantic content, and the Russellian proposition that Hillary Clinton is a politician (i.e. $\langle \text{Hillary Clinton, the property of being a politician} \rangle$) will be the semantic content of every token of 'Hillary Clinton is a politician' that has a semantic content. On the other hand, the expressions 'she' and 'She is a politician' will have non-constant meanings, as, clearly, different tokens of those expression types may have different semantic contents. It is convenient to represent constant and non-constant meanings as partial functions from tokens onto things that may be the semantic contents of those tokens.³⁴ Thus, to a pretty rough first approximation (although not so rough that it doesn't convey the general idea) we might have, for example:

³⁴ Then we should say that two expressions ε and ε' have the same meaning if the functions that are their meanings are the same except that the function ε means has tokens of ε as its domain of arguments, whereas the function ε' means has tokens of ε' as its domain of arguments.

- The meaning of ‘Hillary Clinton is a politician’ = $\langle \tau, f \rangle$, where for any token τ of the sentence, $f(\tau)$ = the proposition that Hillary Clinton is a politician.
- The meaning of ‘Hillary Clinton’ = that function f such that for any token τ of the name, $f(\tau)$ = Hillary Clinton.
- The meaning of ‘politician’ = that function f such that for any token τ of the predicate, $f(\tau)$ = the property of being a politician.
- The meaning of ‘I’ = that function f such that for any token τ of the pronoun, $f(\tau) = x$ iff x = the person who uttered τ & who referred to herself with τ .
- The meaning of ‘she’ = that function f such that for any token τ of ‘she’, $f(\tau) = x$ iff x is the female to whom the speaker referred with τ . τ
- The meaning of ‘She is a politician’ = $\langle \tau, f \rangle$, where for any token τ of the sentence, $f(\tau)$ = the proposition that x is a politician iff for some function g , (i) ‘she’ means g , (ii) τ' is the token of ‘she’ that occurs in τ , & (iii) $x = g(\tau')$.

So much for my brief sketch of Best Bet, the theory that is the best bet for being the semantic component of an internally represented generative grammar if GGH is correct. What Best Bet emphasizes is the foundational importance of the nature of the content attributed to propositional speech acts in the complement clauses that ascribe propositional content to them. For sentence meanings, if there are such things, must be things that constrain both the speech acts sentences may be used to perform and the contents those speech acts may have, and it’s because of the way they constrain those things that speakers are able to mean things in uttering sentences and hearers are able to know what they mean. If GGH is correct, the semantic component of an internally represented generative grammar of a language will pair the expressions of the language with the constraints that are their meanings, and tacit representations of those constraints will be crucially implicated in the information processing that takes one from the

perception of an utterance of a sentence to the knowledge of what the speaker meant in producing it.

IV. *Formal Semantics Again*

The foregoing discussion of Best Bet has implications for formal semantics in linguistics.³⁵ In my initial discussion of formal semantics I claimed that no compositional truth theory—whether or not it’s model-theoretic, and with or without possible worlds—can be the semantic component of a generative grammar that verifies GGH because the theorems of such a theory won’t ascribe properties to sentences that could constrain what could be meant in literal utterances of them well enough to explain how understanding a sentence enables a person to mean something in uttering it or to know what a speaker means in uttering it. That observation will neither surprise nor consternate those linguists, such as Barbara Partee, whose interest in semantic theories is limited to theories that seek only to account for the truth conditions of, and entailment relations among, the sentences of a given natural language. What might concern them is what GGH implies about a model-theoretic possible-worlds truth theory for a language even if the theory’s assignments of possible-worlds truth conditions to the sentences of the language are correct. To keep the points that need to be made from getting lost in complexities, I’ll stipulate *E* to be a variant of English all of whose sentences are unambiguous and truth-evaluable. Now let *T* be a model- and truth-theoretic possible-worlds semantics for *E* whose choice of an “intended” model for *E* entails for each sentence σ of *E* a correct truth condition of the form:

[$T\sigma_{tc}$] For any token σ_τ of σ & any possible world w , σ_τ is
true in w iff ...; false in w iff ...; otherwise
undefined for w .

³⁵ My impression is that as ‘formal semantics’ is currently used in linguistics, any semantic theory that uses symbols counts as falling under the ‘formal semantics’ rubric, even if the theory using the symbols is a Best Bet compositional meaning theory. I trust, however, that by now it’s clear that my discussion of “formal semantics” in linguistics is about the sort of model- and truth-theoretic possible-worlds semantics for natural languages—or generalizations of it, such as dynamic semantics—inspired by Montague’s work.

And let's suppose that GGH is true, and that therefore the semantic component of the internally represented generative grammar for E is a Best Bet compositional meaning theory, M , that, on the basis of its semantic rules and the meanings it assigns in its base axioms to E 's morphemes, generates for each sentence σ a meaning-ascribing theorem

$$\sigma \text{ means } \langle \vdash, f \rangle,$$

where, for some condition C , f is such that, for every token σ_τ of σ & Russellian proposition p , $f(\sigma_\tau) = p$ iff p uniquely satisfies C . Then there will be a true theory of E , M^+ , which, to a first approximation, is simply the conjunction of M with the following propositions (where ' σ ' ranges over sentences of E , ' f ' over functions, ' p ' over Russellian propositions; and ' Φ ' over sentential forms (e.g. sentences of the form ' S or not S '), where the subscript ' m ' signifies that the subscripted modal term expresses a metaphysical modality, while the subscript ' l ' signifies that the subscripted modal term expresses a logical, or formal, modality):

1. For any σ , a token σ_τ of σ is true iff for some f, p , (i) σ means $\langle \vdash, f \rangle$; (ii) $f(\sigma_\tau) = p$; & (iii) p is true. Likewise, *mutatis mutandis*, for the token's falsity conditions.
2. For any x_1, \dots, x_n, R^n , $\langle \langle x_1, \dots, x_n \rangle, R^n \rangle$ is true iff $\langle x_1, \dots, x_n \rangle$ instantiates R^n , false iff it doesn't.
3. For any σ , a token σ_τ of σ is necessarily _{m} true iff for some f, p , (i) σ means $\langle \vdash, f \rangle$; (ii) $f(\sigma_\tau) = p$; & (iii) p is necessarily _{m} true. Likewise, *mutatis mutandis*, for the falsity conditions and for the truth and falsity conditions for a sentence token's being possibly _{m} true/false.
4. For any σ , a token σ_τ of σ is necessarily _{l} true iff σ is necessarily _{m} true & for some form Φ , (i) σ_τ is of form Φ & (ii) if any token of any sentence is of form Φ , then it's necessarily _{m} true. Likewise, *mutatis mutandis*, for the falsity conditions and for the truth and falsity conditions for a sentence token's being possibly _{l} true/false.

(Comments: (1) tells us that a sentence token's truth-value is inherited from that of the proposition that is its semantic content. (2) is what requires me to say that "to a first approximation" the true theory M^+ conjoins M with (1)-(4). (2) is best thought of as what the truth theory for Russellian propositions would be if it weren't for the Liar and other semantic paradoxes, such as the property version of Russell's paradox,³⁶ which suggest that the notion of truth it employs isn't fit for service. This, of course, is also a problem for a model-theoretic approach to natural language semantics, and perhaps the only take-home point that I can make here is that if a suitably regenerated notion of truth for propositions is available and 'true' in (2) is taken to express it, then (2) will hold for the Russellian propositions it's permitted to range over. We may assume that if a kosher notion of truth is available for natural language sentences, then one will be available for propositions, and Best Bet will continue to be a theory that holds that sentence tokens inherit their truth conditions from the propositions they express. (3) imparts an important lesson about modal properties. It's a mistake to assume that a possible-worlds semantics in any sense gives the *meaning* of the modal notions whose logic it aims to capture. How could it? As an account of the meaning of 'possibly', the claim [that a proposition is *possible* just in case there is some *possible* world in which it's true] is blatantly circular. Best Bet would doubtless want to treat the claim *that such-and-such proposition is metaphysically necessary* as ascribing an irreducible property to the proposition. (4) reminds us that logical, or formal, necessity is in the first instance a property of sentences, but still a property that is definable in terms of the modal properties of propositions. We wouldn't want to say that 'Every bachelor is unmarried' is a logical or formal truth even if we wanted to say that it expressed the proposition that everything that is both a man and unmarried is unmarried, but the notion of logical or formal truth is definable in a way that respects the idea that sentences get their truth conditions from the propositions they express.)

³⁶ Some properties, such as the property of being odorless, instantiate themselves (the property of being odorless is odorless), whereas other properties, such as the property of being a dog, don't instantiate themselves (the property of being a dog isn't a dog), so we seem to have the perfectly good proposition *that the property of being a property that doesn't instantiate itself instantiates itself*, yet when that proposition is fed into (2) we get the contradiction that it's false if it's true and true if it's false.

OK, so we have two true theories, *T*, the model-theoretic possible-worlds truth theory for *E*, and *M*, the Best Bet meaning theory for *M*. It should be clear that *M* is a more *fundamental* semantic theory of *E* than *T*. For *M*'s being true *explains* why *T* is true, but *T*'s being true doesn't explain why *M* is true. Moreover, *M*'s being true explains *T*'s being true in a way that demonstrates that there is a way of assigning truth conditions to the sentences of *E* which *supersedes* *T*'s way of assigning them. Barbara Partee says somewhere that while a model-theoretic possible-worlds truth theory for a language can't on its own be a theory of meaning for that language (because a sentence's possible-worlds truth conditions can't be equated with its meaning), nevertheless, those truth conditions will be a large part of the sentence's meaning, and thus the truth theory will be a large part of the language's meaning theory. But even that can't be true: the meaning theory doesn't use the notion of a possible world, and so in no literal sense will a possible-worlds truth theory be "part" of the meaning theory, and, furthermore, the truth theory won't even be *logically* derivable from the meaning theory, since there is no sense in which the meaning theory can be construed as an extension, conservative or otherwise, of the truth theory: a logical entailment is precluded just by virtue of the fact that the truth theory uses notions (e.g. models and possible worlds) that don't occur in the meaning theory. Semantics in linguistics is supposed to be, or so my *Oxford Dictionary of Linguistics* tells me, "the study of meaning." This means that the formal semanticist can't be unconcerned that the kind of semantic theory for a natural language that interests her has no place in a theory of linguistic competence. For if GGH is correct, then the more fundamental semantic theory is the compositional *meaning* theory that is the semantic component of the internally represented generative grammar, and if that is so, then linguistic semantics has so far ignored what really ought to be its primary concern. That of course is consistent with recognizing that there is much to be learned from formal semanticists about the semantics of natural languages, as witnessed, for example, by work done in formal semantics on anaphora and natural language quantifiers.³⁷

³⁷ Versions of this essay were given as talks at the conference of "Language: the Limits of Representation and Understanding" in Erfurt, Germany; SUNY New Paltz; NYU Mind & Language Seminar; and the Semantics and Philosophy in Europe 6 conference in St Petersburg, Russia. I profited from the discussions at each of those talks, but I would like

especially to thank Chris Barker, John Collins, Paul Egré, Giordana Grossi, Hans Kamp, Stephen Neale, Barbara Partee, Paul Pietroski, Zoltán Szabó, and Seth Yalcin.

Works Cited

- Chomsky, N. (1995). "Language and Nature," *Mind* 104: 1-61.
- Cresswell, M. (1978). "Semantic Competence," in F. Guenther and M. Guenther-Reutter (eds), *Meaning and Translation* (Duckworth).
- _____ (1985). *Structured Meanings: the Semantics of Propositional Attitudes* (MIT Press).
- Dennett, D. (1969). *Content and Consciousness* (Routledge & Kegan Paul).
- Field, H. (2008). *Saving Truth from Paradox* (Oxford University Press).
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar* (Blackwell).
- Katz, J. and Fodor, J. (1963). "The Structure of a Semantic Theory," *Language* 39: 170-210.
- Katz, J. and Postal, P. (1964). *An Integrated Theory of Linguistic Descriptions* (MIT Press).
- Lewis, D. (1972). "General Semantics," in D. Davidson and G. Harman (eds.), *Semantics of Natural Language* (D. Reidel Publishing Co.): 169-218.
- _____ (1999a). "New Work for the Theory of Universals," in Lewis (1999b).
- _____ (1999b). *Papers in Metaphysics and Epistemology* (Cambridge University Press).
- Mathews, P. (ed.) (2007). *Oxford Concise Dictionary of Linguistics*, 2nd Edition (Oxford University Press).
- Montague, R. (1970a). "English as Formal Language, in B. Visentini et al. (eds), *Linguaggi nella Società e nella Tecnica* (Edizioni di Comunità): 189-224.
- _____ (1970b). "Universal Grammar," *Theoria* 36: 373-98.
- _____ (1973). "The Proper Treatment of Quantification in Ordinary English," in J. Hintikka, J. Moravcsik, and P. Suppes, *Approaches to Natural Language* (Reidel).
- Partee, B. (1988). "Semantic Facts and Psychological Facts," *Mind & Language* 3: 43-52.
- _____ (2004): "Reflections of a Formal Semanticist," in B. Partee, *Compsiiionality in Formal Semantics: Selected Papers by Barara H. Partee* (Blacwell):

- _____ (2011). "Formal Semantics: Origins, Issues, Early Impact," *The Baltic International Yearbook of Cognition, Logic and Communication* 6: 1-52.
- Schiffer, S. (2003). *The Things We Mean* (Oxford University Press).
- Stich, S. (1978). "Beliefs and Subdoxastic States," *Philosophy of Science* 45: 499-518.
- Tarski, A. (1956). "The Concept of Truth in Formalized Languages," in *Logic, Semantics, Metamathematics*, tr. J. Woodger (Oxford University Press).sa