

Logic and Semantic Analysis

Ernest Lepore and Matthew Stone

July 8, 2005

1 Introduction

When we wish to frame or to communicate a precise and nuanced argument, we should first clarify whatever meaningful distinctions our reasoning exploits. That's why every good paper begins by defining its terms. A tiger is a large and ferocious predatory cat, yellow with black stripes. A bachelor is an unmarried man. Freedom is the capacity to choose one's actions for oneself, independent of causal forces in the outside world. Knowledge is justified true belief. Getting clear on our concepts is the process of analysis. It is such a fundamental part of philosophical practice that the preponderance of contemporary philosophical writing in English today is described as 'analytic'.

Analysis issues in statements that seem to equate ideas. When we agree on these statements, we seem to do so in advance of any argument. When we dispute them, we seem to argue not about what's true or false, but about what's worth discussing at all. These observations naturally lead to a particular interpretation of philosophical practice: When we do analysis, we ask ourselves about the meanings of our sentences or the contents of our thoughts. The meanings and contents we discover are inherent in our sentences and thoughts, and could not have been otherwise. And our introspection comes with the subjective authority with which we know our own minds. So when we practice analysis, the statements we derive are 'analytic'. Traditionally, this has been understood to mean that they are necessarily true, purely in virtue of the meanings of our terms, and that they are knowable *a priori*, independently of any experience.

If there were any analytic statements in *this* sense, it would be a boon for philosophy. We could argue against a skeptic about the real world, for example, by constructing an incontrovertible analysis of our concepts and reasoning from it to derive the knowledge we hope to attribute to ourselves from the incorrigible evidence of our senses. We could also argue against a skeptic about intersubjectivity – someone who contested the idea of public truth or falsehood on the grounds that our ideas are private and even language does not offer public means for us to share them. If we share the terms with which we frame our arguments, and know their meanings *a priori*, we can be confident that we agree on the questions, and so can be confident that our disputes contest the answers.

We shall follow Quine (1953) in thinking that there are no analytic statements in this broad sense. When we do semantic analysis, we cannot hope to infer statements that have a special status in philosophy. We explicate his reasoning in this chapter, and suggest that while developments in the philosophy of mind and language over the more than fifty years since Quine may have superseded much of his perspective, they continue to support his conclusion. However, although we agree that there are no *a priori* truths

that describe linguistic meaning, we do not believe semantic analysis should be abandoned. On the contrary, the enterprise of semantics in linguistics and in the philosophy of language depends on our ability to articulate and formalize precise statements that specify and explicate linguistic meaning. Such statements serve as crucial hypotheses in the scientific enterprise of accounting for linguistic behavior. Only such general laws of meaning can explain the productive capacity that we all have in order to use and to understand the potential infinity of statements in our native language.

Our view reflects the belief that the laws of meaning are not knowable *a priori*. As speakers, we learn them or fix them in our interactions with our linguistic community and the environment; as theorists, we hypothesize them to explain our linguistic competence. Thus, it is not a *logical* consequence of our knowledge of meaning that we share the meanings of our terms with others in our community; it is an assumption that we make but that a skeptic might deny. Moreover, it turns out that the semantic facts we need to explain language use are profoundly boring for philosophy more generally. Here is a typical example: If Kim kissed Sandy, then there was a kissing. Such facts do not bridge sensory and categorical vocabulary. They will not allow us to relate our knowledge of the world to the evidence of our senses; so they will not underpin our epistemology. Analysis, again, has clear limits, and we will have to fight the skeptic the hard way. Yet, all in all, we offer an optimistic view about analysis. If you already believe we have a broadly shared language to discuss an accurately perceived world, then semantic analysis points to a good explanation for how that could be.

Since Quine (1960), and especially since Montague (1974), formalism has become central to the practice of semantic analysis. Logic offers us an appealing, lightweight notation while nevertheless forcing us to be completely explicit and completely consistent in the ontology and concepts we appeal to in theorizing. When our arguments become difficult, this discipline becomes indispensable. The use of logic also anticipates the place for semantic analysis in a cognitive science of language and the world mediated by mental representations. Semantic analysis interprets linguistic meaning in terms of something fundamentally nonlinguistic: relationships in the real world. Here the nonlinguistic metalanguage of formal logic serves as a useful reminder that we have moved outside the realm of grammatical knowledge and communicative convention, into a system that holds its meaning through standards and mechanisms fundamentally different to those of language. In this chapter, then, we will also emphasize how formalism can be exploited (especially first-order logic but also higher-order logic and intensional logic), both to craft particular analyses more precisely and to think more generally about what analysis is.

2 Some classic arguments.

We begin with a review of three philosophical perspectives on knowledge of meaning. First, we present Quine's view of meaning as a theoretical interpretation of a speaker's present dispositions to assent to, or to dissent from, utterances in context. Quine's empiricism helps to make clear the intuition that there are no *a priori* truths about meaning. In fact, Quine argues that a theoretical interpretation is unable in principle to distinguish between facts about meaning and other empirical truths. Then, we discuss a different view of meaning, particularly one associated with Kripke and Putnam. Here the project of semantics is not to interpret a speaker's current behavior but to explain it; this

kind of semantics will show how a speaker's thoughts and utterances can be meaningfully grounded in his prior experience, particularly his interactions with members of his community. This view predicts that there are truths about meaning: we can find statements that, just because of the meanings of the terms they contain, could not have been false. But on this view there still are no *a priori* truths about meaning, for truths about meaning are knowable only *a posteriori*. Finally, we link these considerations to the architecture of contemporary semantic theory, drawing on the interpretive truth-theories of Davidson (1967b/2001), as fleshed out, for example, in Larson and Segal (1995). Linguistic methodology challenges us to explain the systematicity and productivity of meaning in language. One response to this challenge follows Kripke and Putnam, but of course not Quine, in appealing to explicit knowledge of meaning. On this account, we formulate novel sentences and recognize the meanings of others' novel sentences by reasoning directly from facts we have learned about the meanings of expressions of our language – and our mental apparatus stores these semantic facts and computes them with a nonlinguistic formal system – 'a language of thought' (Fodor 1975), for which logic provides an attractive model.

2.1 Quine on analysis

Quine (1960) developed his arguments about analysis as part of a larger project to make sense of commonsense mentalistic descriptions of behavior within a rigorous, scientific framework. Like many philosophers before and since, Quine was deeply skeptical about the constructs and explanations of commonsense psychology. Indeed, in his philosophy, Quine generally regarded all knowledge as incomplete, provisional, and open to revision. To describe an agent's beliefs, Quine dispensed with theoretical constructs such as mental representations and abstract objects such as meanings and propositions. He focused instead on characterizing the sentences which an agent will assent to and the circumstances in which the agent will do so. In this way, Quine sought to refer mentalistic descriptions straightforwardly to observable distinctions; our psychological theories could then exploit these distinctions to interpret the statements and actions of others against our own theoretical picture of the world.

Against this background, Quine argued that meaning is indeterminate. There could be many, equally valid schemes to render the content of another's beliefs in our own terms. Each of these schemes would be holistic, in that it would aim to provide a consistent overall interpretation of a body of beliefs in terms of our own. The key to such an interpretation would be its overall coherence; as far as possible, it would have to interpret true sentences as true and false sentences as false. To judge another to be systematically in error would undercut our claim to understand his sentences. Quine saw this holism of interpretation in mental explanation as continuous with the holism of other kinds of scientific theorizing. For Quine, talk of hypothesized entities in science, like electrons, derives its content from its role in a network of principles that ultimately leads to empirical predictions.

Quine's famous 'Gavagai' thought experiment illustrates both his holistic view of meaning and his empirical, *a posteriori*, understanding of the attribution of semantic content. He invites us to consider a field linguist working with speakers of an isolated language without access to any interpreters. The task of this linguist is to arrive at translations from the speaker's language into his own for heard utterances. Quine calls

this project ‘radical translation’. Of course, the linguist must consult speakers as informants, so she first discovers how they signal agreement and disagreement. Now to get a handle on what a particular utterance means, she determines the situations that prompt informants’ agreement with it, and those that prompt their disagreement with it. She settles on a good translation by finding a statement in her own language for which her judgments match those of her informants. She discovers, for example, that ‘Gavagai!’ can be translated as ‘Look, a rabbit!’ by seeing that informants accept ‘Gavagai!’ in more or less the situations where she would accept ‘Look, a rabbit!’ – in short, when there is a rabbit around.

Can translations so discovered be regarded purely as facts about meaning? Quine argues they cannot. The linguist’s procedure yields translations that inevitably summarize and reflect a broad swath of informants’ empirical knowledge. An informant’s assent to ‘Gavagai!’ will be grounded perceptually, in his own experience of the appearance of rabbits and the contexts where rabbits are found, and it may also be grounded socially, in information provided by his companions or simply in the shared background of his broader culture. In arriving at a translation, the linguist *must* locate a sentence in her own language that approximates the *entire* network of factors which determine her informant’s judgment. She cannot narrow the scope of translation by ruling out some of these factors as extrinsic to meaning. After all, she aims to *discover* her informant’s meanings; if she rules out potential aspects of meaning in advance, she abandons science for mere prejudice. Moreover, in construing some information as extrinsic to meaning, she worsens the fit between her informant’s judgments about his utterances and her judgments about the translations. In this sense, she really does less faithfully capture what her informant means.

For Quine, a child learning her native language is in the same situation as a scientist translating an exotic new language. She must construct a radical interpretation of the utterances of her community. Her conclusions take the form of an overall theory about how people act in her community, and this theory must be assessed holistically against all her empirical evidence. The theory must allow her to understand the sentences members of her community accept as true, and those members of her community reject as false. When she evaluates her interpretation this way, she can make no principled distinction between what ‘rabbit’ *means* in her community, and what rabbits *are* to her community. If members of her community had assented to other sentences involving the word ‘rabbit’, her theory would have done better to interpret ‘rabbit’ differently. This is why Quine suggests that a theoretical interpretation is unable in principle to distinguish between facts about meaning and other empirical truths. No sentences are true solely in virtue of meaning. *A fortiori*, there can be no ‘analytic’ sentences that are true solely in virtue of meaning, that are necessarily true, and that can be known *a priori*.

2.2 Kripke, Putnam, and the necessary *a posteriori*

Quine conceived of interpretation as our effort to *understand* one another. Interpretation is charitable. In constructing an interpretation, we primarily aim to find points of agreement between our perspective and another’s. A different project in semantics is to try to *explain* one another. An explanation gives meaning to another person’s thoughts and sentences by describing his relationship to the world around him. Explanation need not be as charitable as interpretation. Indeed, the appeal of explanatory semantics springs

from thought experiments where we are prepared to judge another to be systematically in error, and yet we refrain from reinterpreting his sentences.

Here is a clear case. With Putnam (1975), imagine a planet exactly like Earth, except that where Earth has water, this other planet, Twin Earth, has another, mysterious substance, XYZ. To human senses, this substance seems exactly the same as water; nevertheless, it has a fundamentally different chemical structure. Imagine further that it's still the year 1700, and chemical structure has yet to be discovered. Still, we judge that the English word 'water', on Earth, means water, whereas the Twin English word 'water', on Twin Earth, means XYZ. Moreover, if an Earthling were suddenly teleported to Twin Earth, he would still speak English, and his word 'water' would still mean water – this despite the fact that he might have exactly the same dispositions as Twin Earthers have to accept or reject statements about his new surroundings. In short, the unfortunate Earthling would think he was surrounded by lots of water, and he would be completely wrong.

You can, perhaps, take Putnam's Twin Earth thought experiments to illustrate why Quine thought translation was indeterminate. When you look at how a speaker is disposed to respond to sentences of English, you can interpret 'water' equally well as water, XYZ, or the disjunction of the two. These interpretations are different, and they would assign different truth values to English statements in meaningful (but ultimately inaccessible) situations. In fact, though, when we say that 'water' in English means water, we are applying a different standard than Quine's. Kripke's (1972/1980) suggestion is that our standard is based on our understanding that English speakers *intend* to pick out a particular kind of stuff in their own environment. As a community, English speakers have encountered this stuff and named it 'water'. And as a community, English speakers work together to ensure first that the community maintains the referential connection between the word 'water' and that stuff, and only secondarily that individuals in the community can themselves recognize examples of the stuff in particular situations. When as observers we recognize that 'water' means water, we aren't summarizing the epistemic abilities of particular speakers. Rather, we are summarizing social commitments and causal connections in the community that have worked across speakers to hook the word 'water' up with the stuff, and keep it that way. What we do, ultimately, is to *explain* how speakers can use language to refer in shared ways to aspects of the world.

Kripke motivates his view with an analogy between words for kinds, such as 'water', and proper names, such as 'Richard Feynman'. In the case of proper names, we can point to the social practices that initially fix the reference of a name and transmit that reference within the community.

Someone, let's say, a baby, is born; his parents call him by a certain name. They talk about him to their friends. Other people meet him. Through various sorts of talk the name is spread from link to link as if by a chain. A speaker who is on the far end of this chain, who has heard about, say, Richard Feynman, in the market place or elsewhere, may be referring to Richard Feynman even though he can't remember from whom he first heard of Feynman or from whom he ever heard of Feynman. He knows that Feynman is a famous physicist. A certain passage of communication reaching ultimately to the man himself does reach the speaker. He is then

referring to Feynman even though he can't identify him uniquely. He doesn't know what a Feynman diagram is, he doesn't know what the Feynman theory of pair production and annihilation is. Not only that: he'd have trouble distinguishing between Gell-Mann and Feynman. (Kripke 1980, p. 91)

The result is that we can judge a speaker's reference with a proper name independent of the sentences that the speaker would assent to or reject.

In the case of common nouns like 'water', the word has had its reference since time immemorial. Nevertheless, new speakers still link themselves into chains of reference that participate in and preserve the connection between 'water' and water. So analogously, we take an English speaker's word 'water' to refer to water, independent of the sentences that the speaker would accept or reject.

Most philosophers find this new view of semantics more satisfying than Quine's in key respects. The new view offers a close fit to our intuitive understanding of ourselves. It seems that we really do commit to use our words with the same reference as our community. And when others make claims about the world, it seems that we really do assess and dispute those claims with respect to the common standard in the community, not just with respect to our interpretation of others' personal theories of the world. Moreover, the new view suggests specific grounds to critique the motivations behind the Quinean interpretation. For example, on this new view, we will inevitably focus on certain aspects of an agent's verbal behavior and not others when we assign meanings to his utterances. We do so because we locate semantics as *part* of a broader science of the mind, which combines a theory of language with a theory of action (including an account of our intentions and social relationships) and a theory of perception (including an account of the limits and failings of our observation). Semantics in itself explains only so much – and, not surprisingly, just because we understand the meaning of someone's sentences, we do not *ipso facto* understand *them*.

Crucially, this new view of semantics predicts that some statements are necessarily true, solely in virtue of the meanings of the words involved. We've already seen that it is a fact about meaning that 'Richard Feynman' names Richard Feynman, or that 'water' names water. We can go further. 'Hesperus' names the planet Venus, 'Phosphorus' names the planet Venus, 'is' names the identity relation. So sentence (1) follows, just as a matter of meaning:

1. Hesperus is Phosphorus.

Given that 'Hesperus' and 'Phosphorus' are both names for the planet Venus, (1) has to be true. There is no way that *that planet* could have failed to be *that planet*. Like sentence (1), the other facts that follow from the meanings of our language are necessarily true.

However, these facts about meaning are not knowable *a priori*. We discovered them. To illustrate the point, imagine that, early on, the Ancient Greeks were in an epistemic situation that left it open whether the bright object that sometimes appeared in the morning sky was the same as the bright thing that sometimes appeared in the evening sky. They could not distinguish themselves from their doubles on a Twin Earth where the morning star and the evening star actually were distinct objects (alien satellites, we might suppose). These Twin Earthers would speak a language in which (1) translates

into a false sentence – indeed, a necessarily false sentence. For the Ancient Greeks, however, the translation of (1) was necessarily true. Eventually, the Ancient Greeks advanced their science, and improved their epistemic situation. They realized that in our case there is only one celestial object. At the same time, then, they discovered that (1) is necessarily true.

When we reflect on the generality of the Twin-Earth thought experiments, it is clear that semantic facts are knowable *a posteriori*. We can imagine being quite wrong about what our world is like. In these imaginary situations, our empirical errors extend to errors we make about what our words mean. And, of course, we can also imagine disagreeing with others about what the world is like. Though we are committed to use our words with the shared reference of our community, we must be prepared to resolve our dispute by relinquishing facts that we think are necessarily true – facts we think characterize the meanings of our words and the contents of our thoughts. Thus, the situation for philosophy generally remains much as Quine understood it.

Contra Quine, we can pursue semantic analysis, and attempt to discover statements that are necessarily true in virtue of the meanings of the terms. But any insights we derive about meaning are *a posteriori*. So, as Quine concluded, they are hypotheses that we make and test like other empirical statements; they do not have a special status in our arguments in virtue of the fact that they characterize meaning.

2.3 Knowledge of meaning and linguistic competence

The development of linguistic semantics as a subfield of cognitive science has confronted philosophical views of meaning with new evidence of the richness and sophistication of speakers' semantic judgments about their native language. To start, speakers are able to understand and produce indefinitely many sentences – sentences neither they nor anyone else in their community has uttered before. This shows their knowledge of language must be *productive*; it must go beyond a fixed lexicon of predefined static elements, and must include a generative system that actively composes linguistic knowledge so as to describe arbitrarily complex structures. The hallmark of productivity in language is recursion. Recursive patterns of complementation, as in (2), and recursive patterns of modification, as in (3) and (4), allow phrases to be nested indefinitely many times within a single sentence.

2. Chris thinks that Kim thought that Robin wanted Sandy to leave.
3. Chris bought a gorgeous new French 3-quart covered copper saucepan.
4. Chris is writing a book that describes inventors that have built machines that changed the world that we live in.

Speakers' capacity to formulate and recognize an open-ended array of possible sentences shows how acute a problem it is to coordinate meaning across speakers. When we learn the meaning of our native language, we must generalize from the finite record of our previous experience to an infinity of other sentences and situations. If we thereby arrive at a common understanding of the meanings of sentences, it must be because language is structured by substantive and inherent constraints that we are able to exploit. The same goes for theoretical projects in semantic analysis. We should expect that our discoveries about the meanings of sentences will generalize to an indefinite array of cases, and that

they will respect and illuminate linguistic constraints. More generally, if our discoveries in semantics are to help explain how speakers can use language meaningfully, we should expect that the generative mechanisms we postulate as theorists will be compatible with the psychological mechanisms that underlie speakers' abilities.

A straightforward way to explain productivity in semantics is to assume that the meanings of particular sentences can be calculated by inference from general facts about meaning in the language. For example, consider the theory presented in examples (5)-(7).

5. 'snow' is a noun phrase and refers to the stuff snow.
6. 'white' is an adjective phrase and refers to the property whiteness.
7. If N is a noun phrase and refers to the stuff S and A is an adjective phrase and refers to the property P, then $\lceil N \text{ is } A \rceil$ is a sentence and is true if and only if S is P.

From this theory, we can derive (8) as a logical consequence.

8. 'Snow is white' is true if and only if snow is white.

We can take (8) as a characterization of meaning, because it links up a sentence, in this case the sentence 'Snow is white', with a condition in the world stated in objective terms, in this case, the condition that snow is white. As semantic theorists, we can use this kind of theory, which Davidson (1967b/2001) calls an interpretive truth-theory, to provide a general account of how sentences link up with conditions in the world. We use atomic formulas to axiomatize the semantics for the elementary structures in the language and use conditional formulas to describe the semantics of complex structures in the language as a function of the semantics of their constituents. We then reason logically from the axioms to associate particular sentences with conditions in the world. As in (5)-(7), this inference will be *compositional*, in that the conclusions we derive will be inferred through a logical derivation that mirrors the syntactic derivation of the sentence.

There are two ways to view interpretive truth-theories such as (5)-(7). We can use an interpretive truth-theory to formulate a semantic theory for a new language. For example, we could be pursuing radical translation. In this case, we are interested in systematically articulating translations for sentences in the object language in terms of sentences in our own; we understand these translations to be derived by inference from the axioms of the theory. This is compatible with a holistic, interpretive view of meaning like Quine's – indeed, Davidson himself subscribed to such a view.

Another way to view interpretive truth-theories such as (5)-(7) is as ingredients of speakers' psychology. On this view, we regard axioms such as (5)-(7) as generalizations that native speakers know tacitly about their language. When speakers formulate or recognize particular utterances, they reason tacitly from this implicit theory to derive conclusions such as (8), about specific new sentences. On this understanding, interpretive truth-theories offer an explanation of how speakers' knowledge of meaning and their inferential abilities together underlie their semantic competence.

This psychological understanding of interpretive truth-theories meshes better with an explanatory view about truth about meaning than with Quine's conceptual framework. For this psychological understanding commits us to a *representational theory of mind*, in which psychological processes involve the tokening of symbolic structures and these

structures are locked on to content about the world through causal or informational mechanisms. (For more, see Fodor, 1975 or 1987.) To see this, let us return to (8):

8, 'snow is white' is true if and only if snow is white.

As cognitive scientists, we expect that speaker knowledge of linguistic meaning will place words and sentences in correspondence with represented regularities in the world. In this case, the correspondence links the sentence 'snow is white' with an answer as to whether snow is in fact white: a distinction you have a handle on some other way. Theories of mental content that postulate causally-grounded nonlinguistic structured representations give you a way to make sense of this idea. A speaker's mental representation of the distinction whether snow is white is composed from mental symbols SNOW for the stuff snow and WHITE for the property of whiteness. The meaning of these symbols follows from their intrinsic role both in the speaker's mental life and in his interactions with his environment. Further symbols, again with a suitable interpretation, represent the property of truth (TRUE) and the logical biconditional (\leftrightarrow). On a representational theory of mind, these symbols can now be assembled into representation (9) corresponding in form and meaning to (8):

9. TRUE('snow is white') \leftrightarrow WHITE(SNOW)

For an individual speaker, the tacit derivation of this symbolic structure can fix the meaning of the sentence 'snow is white'. That is, (9) represents a hypothesis about the meanings of sentences in the speaker's language that the speaker is committed to. At the same time, it shows how the speaker's sensitivity to the meaningful connection between language and the world can be founded directly on the natural intentionality of the speaker's concepts and thoughts.

The willingness to appeal to the representational theory of mind as part of a scientific reconstruction of commonsense psychological explanation is one of the signal departures from Quine that motivates analysis in contemporary philosophy. On Quine's behavioristic perspective, we capture mental content directly through a network of beliefs, circumstances and behaviors. We must assign meaning to agent's sentences without recourse to intermediate representations. And that means statements such as (9) are not special; they are known to be true only as part of a holistic explanation of all the empirical evidence – a speaker does not know that they are true in virtue of meaning. For Quine, (9) can summarize the meaning of 'snow is white' but it cannot establish it. In this sense, the representational theory of mind is essential if we want to explain speakers' knowledge of meaning by appeal to their tacit knowledge of a theory of truth.

3 Doing semantic analysis by analogy to logic

In §2, we sketched some arguments for the position that semantic analysis is both possible and philosophically important in characterizing speakers' rule-governed knowledge of linguistic meaning. This view invites an analogy between the semantics of natural languages and the semantics of the artificial languages of formal logic. The analogy goes back to Frege (1879/1967), who took logic to clarify the features of natural language essential for correct mathematical thought and communication. Davidson's (1967b/2001) interpretive truth-theories, exhibited in (5)-(7), are built on Tarski's (1935/1956) ideas about the semantics of formal languages. The work of Montague in the 1960s and 1970s took the analogy further (Montague 1974). Montague explicitly

advocated an exact parallel between the semantic analysis of English — what ordinary speakers actually know about their language — and the semantics of intensional higher-order logic. In fact, many techniques originally developed for giving semantics to logical languages turn out to be extremely useful in carrying out semantic analysis.

We can briefly describe two notable examples that illustrate this. The first comes from the semantics of first-order logic, where Tarski (1935/1956) introduced assignment functions to give semantics to occurrences of variables. The problem is to explain the truth conditions of complex logical formulas such as (10) in terms of their constituent logical formulas, such as (11). In effect, we have to answer what x means.

10. $\forall x Fx$.

11. Fx .

Of course ‘ x ’ is a variable — it means whatever we want it to mean. Our semantics should simply make this precise. One way to do this is to use a *function* from the set of variables to the set of values—this function then represents a specific set of choices about what the variables correspond to. Such functions are called assignment functions; we use g to name a representative assignment function. Thus (11) is not simply true or false; it is true or false with respect to a set of choices about the variables. Specifically, (11) is true with respect to an assignment g just in case the individual picked out by $g(x)$ has the property represented by F . This framework allows us to understand (10) intuitively in terms of (11): if (11) comes out true for every choice we make for the value of x , then (10) is also true. More precisely, (10) is true with respect to an assignment g just in case (11) is true for any assignment g' exactly like g except possibly that $g'(x)$ differs from $g(x)$.

This treatment of variables and quantifiers provides a useful tool for approaching the analogous meanings of corresponding constructions in natural language, such as quantified noun phrases, bound pronouns and traces. For instance, consider sentence (12). Linguists understand (12) as composed through a syntactic derivation that makes the scope of the quantifier explicit but still indicates that the quantifier supplies the reference for the subject of the sentence and for the pronoun *her*.

12. Every student finished her homework.

The logical form (or LF) in (13) represents these relationships by one common convention: the quantifier is *indexed*, or associated with a variable i ; it is *moved*, so that it brackets a sentence that indicates its scope (just like a logical quantifier); it *leaves a trace*, written as t_i , a variable in the subject position of *finished*; and the pronoun *her* is also indexed with the same variable i .

13. Every student _{i} (t_i finished her _{i} homework).

14. t_i finished her _{i} homework

This LF representation invites us to specify the semantics of (13) in terms of the semantics of (14), and the formalism of assignment functions originally developed for logical languages works beautifully for this. (13) and (14) are related just as (10) and (11) are. For a philosophically-accessible tutorial of the use of assignment functions and other logical techniques to interpret natural language quantifiers, see Pietroski (2006). Analogies between pronouns and variables in formal languages continue to provide

fertile grounds for semantic exploration. For example, dynamic semantics is an approach to the analysis of pronouns in natural language that manipulates assignment functions using formal tools originally developed to account for the meanings of variables in the formal languages used to specify computer programs; see for example Groenendijk and Stokhof (1990) or Muskens (1996).

A second arena where techniques originally developed in the semantics of formal languages provide instructive tools for natural language is in the treatment of time and possibility. Temporal and modal logics feature sentential operators that say when or in what hypothetical conditions a sentence is true. For example, in modal logic, the formula $\diamond p$ says that p is possibly true. Here the semantic problem is to specify how a formula that says that p is possibly true can get its semantics in terms of a constituent formula p , which says that p is actually true. Again, we need to be more precise about the contribution that constituent formulas make to the truth conditions of larger formulas. In $\diamond p$, we don't use p to say that the proposition is actually true; we only say that the proposition is true in a hypothetical possibility that is under consideration. We need to relativize our formal definition of truth accordingly.

Suppose then that we have a set of possible *points of evaluation*, corresponding to the different times, possible worlds, or other situations our logic describes. We use w as a variable ranging over these points of evaluation. The semantics for the logic will then describe whether a formula is true at a particular point of evaluation. Thus, for example, $\diamond p$ is true at a point w if there is some other point w' that is possible from w and p is true at w' . Normally, we distinguish one of the possible points of evaluation as the present one. A formula is simply true if it is true at this present point. See Fitting and Mendelsohn (1998) for a philosophical presentation of the resulting logics.

Of course, natural languages have constructions that seem to depend on times and possibilities in ways analogous to the formal languages. We list a couple in (15) and (16).

15. It's possible that Sandy left.

16. Whenever it rains, Chris brings an umbrella.

To describe the contribution 'Sandy left' makes to (15), or 'Chris brings an umbrella' makes to (16), it is natural to appeal to a notion of truth at a point of evaluation that includes temporal and modal information. (15) is true at a point w if there is some point possible from w where 'Sandy left' is true. (16) is true at a point w if every temporal point comparable to w where 'it rains' is true is also one where 'Chris brings an umbrella' is true.

As in the case of pronouns, pursuing this analogy in detail leads to a rich and profitable interplay between logic and linguistics. When we explore natural language semantics systematically, we find that the tenses of natural languages have some of the properties of operators in temporal logic, but they also have some of the properties of pronouns (Partee 1975; Hinrichs 1985; Webber 1986). Modality in natural language is similar (Roberts 1989). Logically, this behavior finds a match in hybrid logics (Areces *et. al.* 2001) — formal extensions of modal logic that allow formulas to mark the points visited in evaluating a formula and to return to those points elsewhere to describe them further.

While progress in formal semantics for natural language has often been supported by parallels between language and logic, it remains true that philosophical argument often requires us to appreciate the differences between everyday language and idealized uses of formulas in familiar logics. Language depends on context. Logic, by contrast, is so useful in part because it doesn't depend on context: we use logic to present our reasoning in a form whose correctness is intrinsic to its explicit structure and can be determined once and for all. As we explain in more detail in §§3.1-3.2 formal methods for the treatment of context-sensitivity in natural language complicate the picture familiar from logic. Accordingly, we need more nuanced distinctions to describe semantic relationships in languages with context-sensitivity than in languages without it.

Another important distinction is that logical formulas are regimented with specified semantics so that they wear their meanings on their sleeves. By contrast, as we emphasized in §2, natural languages must be learned. We can be uncertain about the meanings of sentences, and we can be creative in how we use them. As we explain in more detail in §3.2, formal accounts of the interpretation of natural language must take this flexibility into account. Again, we find that we need more nuanced distinctions in natural language to describe the relationship between the meaning of a sentence and the way we understand it than logical languages typically call for.

3.1 *Indexicals, demonstratives, and context sensitivity*¹

English has two kinds of terms that are not typically included in the languages of formal logic: demonstratives and indexicals. Demonstratives are words like 'that', 'there', or 'he'. Paradigms of indexicals are 'I', 'now', and 'here' (at least on one of its senses²). What a token of any of these terms refers to depends crucially on the context in which it is tokened. A token of a demonstrative word requires an associated *demonstration* to fix its referent. Typically, but not always, a demonstration is a visual presentation of a salient local object discriminated by a pointing gesture. The token of a demonstrative expression refers to whatever its accompanying demonstration demonstrates – what we shall call a *demonstratum*. Unlike a demonstrative, a token of an indexical requires no accompanying demonstration to secure its reference. The reference of an indexical is determined directly by information about the situation in which it is tokened.

Context sensitivity is a challenge to the analogy between the natural language and logic. As we shall see, context sensitivity requires us to depart from the formalism of logic; this is neither a surprise nor, in itself, a worry. More significantly, context sensitivity requires us to make distinctions in analyzing utterances that we do not need when we describe logical statements. These new distinctions are significant. Chief among them is the distinction between *linguistic meaning* and *utterance content*.

The *linguistic meaning* of an expression is what competent language users learn when they learn that expression. Learning the linguistic meaning of an indexical expression is learning how to adjust its reference across contexts of use. The linguistic meaning of 'tomorrow' determines what that word picks out across varying contexts of its use, as in the designation rule (T).

¹ This section is largely a summary of the important work of Kaplan (1989).

² 'here' is an indexical in 'I am in here', but a demonstrative in 'In two weeks I will be here' (pointing at a city on a map).

(T) In each context of use, ‘tomorrow’ refers to the day after the day of that context of use.

The same goes for demonstratives. Anyone who comprehends the linguistic meaning of a demonstrative expression knows that a satisfactory use of it requires the speaker to make an accompanying demonstration, so, e.g., anyone who understands the demonstrative word ‘that’ must know its linguistic designation rule (D).

(D) A token of ‘that’ refers to the object demonstrated by whoever uses that token.

Linguistic meaning underwrites translation. Suppose that Mike utters (17) yesterday:

17. The Yankees will win tomorrow.

(18) is the translation of (17) into Italian. The reference of ‘domani’ is fixed in Italian by a rule like (T) analogously to ‘tomorrow’ in English.

18. Gli Yankees vinceranno domani.

If Massimo utters (18) on the same day that Mike utters (17) then the two have said the same thing. Thus, in one sense, two sentences with the same linguistic meaning say the same thing.

However, there is another sense in which two sentences with the same linguistic meaning can say different things. *What you express* in this sense is the *content* of your utterance. Distinct utterances of ‘That’s a flower’ can disagree in truth-value because the objects demonstrated might be different. If your two tokens of ‘that’ demonstrate distinct objects, obviously you were not discussing the same object on both occasions. But if your topic of conversation changes, then what you say (or express) on these distinct occasions cannot remain the same. This might be so even though the distinct circumstances surrounding your two tokens might go unnoticed by you (and your audience). The lighting in the room, the sorts of gestures that you make, and every other aspect of your accompanying demonstrations might be indistinguishable to both you and your audience. Demonstratives are quite special linguistic devices in this way: we can token the same words with the same linguistic meaning in contexts that are indistinguishable to the speaker and auditor, and yet, unwittingly, say (or express) different things because different things are referenced.

We can also use different words with *distinct linguistic meanings* to say (or express) the same thing. Should we say, ‘He is a very nice person’ pointing at Bill, and you speaking to Bill say, ‘You are a very nice person’, then we agree in what we say about Bill, even though our words ‘you’ and ‘he’ differ in linguistic meaning. In sum, the same context-sensitive expression can be used on different occasions to say something different, and different context-sensitive expressions distinct in linguistic meaning can be used on the same (or different) occasions to say the same thing. These are the sorts of considerations recommend that we distinguish two sorts of meaning contributions demonstrative expressions can make.

This distinction has pervasive consequences for semantic analysis. We have to cash the distinction out formally, by specifying the meaning of utterance in terms of an occasion of use distinct from its circumstances of evaluation. And we have to follow the distinction conceptually, particularly when we think about whether what we say is *a priori* or metaphysically necessary.

Earlier we observed the usefulness of describing the temporal and modal interpretation of utterances by analogy to the interpretations of modal and temporal logics. Sentence (19) is a sentence in future tense without an (explicit) indexical noun.³

19. The earth *will be destroyed* by a large asteroid.

An utterance of (19) is true at a time t just in case at some later time t' sentence (20), which is in present tense, is true.

20. The earth *is destroyed* by a large asteroid.

Sentence (6) is also in future tense, but unlike (19), (21) has in it an explicit context sensitive expression, namely, the word 'now'.

21. Everything *now* red will be blue.

(or, 'It will be the case that everything *now* red is blue'.)

Were we to reason by analogy with (19), we would have to conclude that (21) is true at time t just in case its present tense counter-part (22) is true at some time t' later than t .

22. Everything now red is blue.

But that's impossible, since (22) is true at a time t' just in case everything that is red *at time t'* is also blue at t' , and such circumstances are not physically possible. Now since an utterance of (21) can be true, it follows that the truth of (21) at a time t cannot require the truth of (22) at a time t' , later than t .

Why does the analogy between (19) and (21) break down? It does on account of the indexical expression 'now' in (21). In evaluating an utterance of (21), we simply cannot ignore that the word 'now', by virtue of its linguistic meaning alone, must pick out the time at which (21) is tokened (used, uttered): so (21) is true at a particular time t just in case (23), which lacks any context sensitive noun, is true at some time t' later than t ,

23. Everything (that is) red *at t* is blue.

where, let's assume, ' t ' is a proper name that refers to a particular time t , namely, the time at which (21) was uttered. And these conditions seem to get things just right. An utterance of (21) is true at a time t just in case (23) is true at some later time t' , i.e., just in case those red objects at t have changed color to blue at t' . The indexical expression 'now' tokened at t , by virtue of its linguistic meaning alone, must pick out the time of its utterance, in this case, t .

When evaluating an utterance of a sentence in future tense, obviously, the relevant circumstances for its evaluation (those circumstances that the utterance is about) are in the future. (21), by virtue of its future tense, is about circumstances after any of its utterances. Likewise, when evaluating a sentence in past tense, the relevant circumstances of evaluation for any of its utterances (those circumstances that utterance is about) are in the past; and so on.

What makes indexical expressions special is that the references of their tokens cling to the occasions on which they are tokened, and these occasions need not be the relevant circumstances for evaluating the truth or falsity of an utterance of sentences in which the indexicals occur. So even though a context in which someone utters (21)

³ We say 'no explicit', since tense itself introduces an indexical element.

cannot by itself determine whether that utterance is true, that context and that context alone, and not, e.g., any future circumstances, suffices to determine the referent of its indexical expression 'now'.

Turning from tense to modality, the same sort of disanalogy arises for sentences about *non-actual* circumstances. Consider the modal sentence (24) with its context sensitive noun 'I'.

24. It is possible that *I* do not exist.
(or, 'I might not have existed'.)

An utterance of (24) does not claim that its non-modal counterpart (25) can *possibly* be uttered truthfully.

25. I do not exist.

(25) can never be uttered truthfully. The linguistic meaning of 'I' is given by the designation rule (R).

- (R) In each context of use, a token of 'I' refers to the agent of that context.

So (25) is false if 'I' has a referent in that circumstance, because the speaker exists, and it lacks any truth value otherwise, because if 'I' fails to secure a referent. Therefore, if an utterance of (24) is true only if its non-modal counterpart (25) is true in a possible circumstance, then it follows that every utterance of (24) is false.⁴ But our intuitions tell us otherwise; every utterance of (24) is true (assuming no necessary beings exist).

Again, we must distinguish between occasion of use and point of evaluation. For any utterance of (24), its token of 'I' picks out *the speaker of that utterance*. So, in evaluating an utterance of (24), it doesn't matter whom the word 'I' can pick out in a possible circumstance. Instead, we fix as the referent of a token of 'I' the actual speaker, and, and ask of this actual speaker whether she exists in some possible circumstance of evaluation. So, if the circumstances are ones in which the speaker's parents failed to reproduce; or circumstances in which the universe came to an end before the speaker was born, then in those circumstances the speaker in question obviously does not exist. Since such circumstances fail to include this speaker, it follows that, unless the speaker is a necessary existent, (24) is true in every context of utterance.

Understanding context sensitivity and its consequences for semantic analysis gives technical insight to philosophical problems. For example, they show how we can learn from statements that are metaphysically necessary. Suppose that we are indirectly watching someone moving about a room on a live video while, at the same time, we are unwittingly watching him directly moving about in front of us. After a lapse of time, we come to notice the strong similarity of the two individual's actions, and we might thereupon begin to wonder whether that guy [pointing to the man in the video] is the same as that guy [pointing directly at the same man]. After a while, let's suppose, we come to assert (26).

26. He is identical to him.

⁴Or, at least not true, since in some circumstances no agent exists, as at the center of the sun or on earth before life formed.

What may not yet be clear to you is that (26), indexed to the context so described, expresses a necessary truth. There are no possible circumstances under which what is said by the utterance of (26) so described might be false. Yet isn't it obvious that something informative was said with this utterance of (26), since we can wonder about truth of the utterance without incurring a charge of linguistic incompetence? This means that whatever an utterance of (26) expresses, it cannot be that we know it to be true *a priori*. So, we see how necessity and *apriority* can come apart. We return to such cases in §3.2 below.

But first we illustrate the complementary point that necessity and *apriority* can come apart the other way. Here is an example of *a priori* knowledge of a contingent truth, which has a Cartesian flavor. Anyone who understands sentence (27) *knows* that no utterance of it can be false.

27. I am here now.

Yet no utterance of this sentence or 'I exist' is *necessarily* true. Clearly, any one of us might have been somewhere other than where he currently is; or, he might have been there at a different time; or he might have been replaced there at that time by someone else, or he might not have existed at all. However, a speaker's knowledge that every utterance of (27) is true is *a priori* knowledge. Contrast (27) with sentence (28). (28) is true under exactly the same conditions as Lepore's utterance of (27).

28. Ernie Lepore is in New Brunswick at 10 a.m. on Thursday, January 27th, 2000.

His knowledge or anyone else's of it is not *a priori*. Internal reflection alone can't help him decide whether he is Ernie Lepore. He might be deluded. Nor can it help me determine whether his current location is in New Brunswick or what date it currently is.⁵

3.2 *What sentences mean and what we understand*

One of the goals of semantic analysis, as we have seen, is to show how speakers' knowledge of natural language allows them to use utterances meaningfully. Thus far, we have addressed this point only in a very preliminary way: if speakers know recursive truth conditions for their language, then they can use these rules to calculate the truth conditions for arbitrarily complex new sentences. But even if speakers can infer the new truth conditions, how exactly does this knowledge help in communication?

A preliminary, approximate answer can be found in the analogy between language and logic. We begin by idealizing the information mutually available to the interlocutors in a conversation as their common ground (Stalnaker, 1973). The common ground settles questions whose answers are uncontroversial, in that interlocutors know the answers, know they know the answers, and so forth. Meanwhile, the common ground leaves open a set of possibilities about which there is not yet agreement: maybe there is a matter of fact that could turn out (for all the interlocutors know) to be one of various ways, or

⁵The more puzzling aspects of this sort of *a priori* knowledge might be eased if one thinks about the worlds in which these sorts of indexical statements would be false. As Descartes convincingly argued, the minimal experience required in order to consider whether these statements are true will require that you know them to be true. Those circumstances in which these sorts of statements are false are circumstances in which someone obviously cannot form the beliefs in question since he does not exist.

maybe the interlocutors actually do know how it turns out but do not realize the knowledge is shared — so it could be that the others know, and it could be that they don't — and so forth. Drawing on the model theory of modal logic, we might represent the possibilities in the common ground as a set of possible worlds. The model theory of modal logic likewise offers a set of possible worlds to represent the proposition associated with any formula. These are the worlds where the formula is true.

If we adopt this logical picture, we can formalize the effect that asserting a formula has on the common ground. When interlocutor *A* asserts a formula *f*, *A* introduces into the conversation the information that *f* is true. Beforehand the common ground is some set of worlds *C*. Afterwards the common ground also takes into account *f*. This formula *f* restricts the live possibilities by requiring the worlds in the common ground to make true a further proposition *p*. So the change that happens when *A* asserts *f* is that the common ground goes from *C* to $C \cap p$. This concise model forms the basis of a range of research characterizing the relationship between truth-conditional semantics and conversational pragmatics in formal terms; see Beaver (2001) for a recent overview.

This idealization has limits. And it's easy to come up with strange puzzles when one moves (perhaps inadvertently) beyond the limits of these idealizations. Here is such a puzzle. Why is it that a speaker would ever assert an identity statement like 'Hesperus is Phosphorus', sentence (1)? The trigger for this puzzle comes from Kripke's arguments, reviewed in §2.2, that sentence (1) is analytic: it is true in all possible worlds. If that is the case, consider what happens when *A* asserts (1). We update the common ground *C* by intersecting it with the set of all possible worlds—the proposition expressed by 'Hesperus is Phosphorus'—leaving the same set *C*. *A* has done nothing; the interlocutors' information has not changed at all! This result is absurd. What has gone wrong?

In fact, in assuming that assertions update the context with the proposition they express, we have implicitly assumed that the participants in the conversation have certain and complete knowledge of their language. For example, interlocutors can calculate that 'Hesperus is Phosphorus' expresses a necessarily true proposition only if they can calculate that 'Hesperus' names Venus and 'Phosphorus' names Venus. Of course, under such circumstances, they don't learn anything from the sentence. But it is also easy to see how this assumption could go unnoticed in extrapolating from logic to language. When we do logic, we are typically interested in inferential relationships that hold no matter what our terms refer to. We assume that the reference of our terms—precisely what matters in the 'Hesperus is Phosphorus' case—is not at issue.

Let us consider how to formalize uses of sentences in more realistic situations. As we do so, we have to be very careful to respect the intuitions of Kripke and Putnam's thought experiments (Stalnaker 1978). Suppose interlocutor *B* does not know that Hesperus is Phosphorus. Recall from our discussion of (1) in §2.2 that what that really means is that *B* cannot distinguish between two possible situations. In the first, there is only one heavenly body out there, and *B* and *B*'s community speak a language English₁ where both *Hesperus* and *Phosphorus* are rigid designators for that body. In the second, there are two distinct heavenly bodies, and *B* speaks a language English₂ where *Hesperus* is a rigid designator for one of them and *Phosphorus* is a rigid designator for the other. Since these possibilities are open for *B*, they must both also be represented in the common ground.

Now, we need a correspondingly expressive notion of assertion. When interlocutor *A* says something, *A* is committed that it is true according to the standards for reference that prevail in the community. Any assertion that *A* makes should turn out to be true in the language *A* speaks. What we have just seen is that any point of evaluation *w* in the common ground could potentially have its own language English_{*w*} with relevant differences from English as spoken in the real world. Adapting Stalnaker's (1978) terminology,⁶ we can associate any utterance *u* with a *diagonal proposition*: this proposition is true at a point *w* if the proposition that *u* expresses in English_{*w*} is true in *w*. In the case of *Hesperus is Phosphorus*, the effect of *A*'s assertion is to intersect the common ground with this diagonal proposition. Concretely, we retain in the common ground worlds of the first kind, where English₁ is spoken, Hesperus and Phosphorus are necessarily the same and *A*'s assertion is necessarily true. However, we discard from the common ground worlds of the second kind, where English₂ is spoken, Hesperus and Phosphorus are necessarily different and *A*'s assertion is necessarily false.

There is substantially more to be said about the relationship between utterance meaning and the information that interlocutors convey. Speakers are rational, so they can communicate indirectly through the choices they make about what information to give and how to present it (Grice, 1975). And language use is collaborative, so speakers can often count on hearers to accommodate information they implicitly presuppose (Lewis, 1979; Thomason 1990) and to recognize projects they indirectly suggest (Searle, 1975). The element of strategy continues to inform fine-grained accounts of linguistic phenomena at the boundary between semantics and pragmatics — for everything from the resolution of pronouns (Beaver 2004) to the use of vague language (Kyburg and Morreau, 2000; Barker 2002). Throughout this research, as in the example from Stalnaker presented above, the need to consider the knowledge and choices of actual speakers remains a central theoretical construct that distinguishes semantic analysis from logic.

4 Doing semantic analysis in logic

Logic offers more than an analogy for doing semantics for natural language. Logic is a tool that makes semantic analysis easier to do, to present and to understand. It allows us to be precise about the ontology that our semantics presupposes — what knowledge of different kinds of entities and relationships best explains our ability to use language. It allows us to hook up that ontology precisely with the syntax. And it allows us to use the mathematics of logical consequence to describe the semantic relationships among

⁶ Stalnaker's presentation of the idea is rather more complicated. Stalnaker's idea is to use possible worlds as representations of the context of an utterance, and to treat the reference of proper names as a kind of context-sensitivity. We think it's better to distinguish between interlocutors' epistemic context, which describes the information they share, and their linguistic context, which describes the state of their interaction. For example, the interaction could have a well-defined state that only one of the participants knows about; this would be part of the linguistic context but not the epistemic context. Analogously, we think it's better to distinguish between interlocutors' uncertainty *about* their language and the context-sensitivity of terms *within* their language. Stalnaker concludes that it is a kind of Gricean repair when an utterance conveys the diagonal proposition in his sense of factoring out all kinds of context-dependency; but we see no reason why an utterance should not always convey the diagonal proposition in our sense of factoring out dependence on the language spoken.

sentences. This mathematics shows how our judgments might be grounded in computations over representations of knowledge of meaning. In §4.1, we describe one of the most influential theories of the sort, Davidson's (1967a/2002) use of events to specify the logical form of verbs of change. Such philosophical analyses are important to semantics, but, as we emphasize in §4.2, these analyses are most successful when they are most banal. We are generally skeptical that semantic analyses that explain how word meanings compose to yield sentence meanings will lead to substantive insights about the internal structure of word meanings themselves. Thus, the project of semantic analysis needs to be clearly distinguished from the broader project of using logical formalisms to model mental representations and cognitive processes in general.

4.1 *Verbs of change and their modifiers*⁷

Verbs of change are verbs that correspond to events, processes or actions which occur, go on, or take place. For example, the verbs 'to leave' and 'to kiss' are verbs of change. Such verbs can be modified by prepositional phrases, as in 'John left *for* the store' or 'Frank kissed Mary *in* Tompkins Square'. This sort of modification seems productive, since there is no obvious limit to how many modifying prepositional phrases we can modify a verb of change with, as the sequence of sentences (29) through (33) suggests.

29. Mary kissed John.

30. Mary kissed John in Tompkins Square.

31. Mary kissed John in Tompkins Square after midnight.

32. Mary kissed John in Tompkins Square after midnight behind his left ear.

33. Mary kissed John in Tompkins Square after midnight behind his left ear on August 24, 1999.

If this process of expansion can be continued indefinitely, then English can express an open-ended set of propositions just by virtue of adding prepositional phrases to a sentence that modify the verb.

Interpreting (29)-(33) thus requires reaching into our bag of learned expressions in order to compose their meanings into meaningful complexes. One promising proposal for how to do this is provided by what we shall call the *event analysis*. According to this proposal, (29) and (30) are best understood to be synonymous with (34) and (35).

34. *There is an event that was a kissing by Mary of John.*

35. *There exists an event that was a kissing by Mary of John and that was in Tompkins Square.*

That is, events figure explicitly in the rules and computations English speakers use to formulate and understand sentences like (29) and (30).

This is a substantive discovery about the meaning of English. Given their surface form, you might have thought that (29) involved a two-place predicate 'kissed' relating Mary and John, and you might have thought that (30) involved a three-place predicate 'kissed in' relating Mary, John and the park. However, the analyses in (34) and (35) urge

⁷ This section is largely a summary of Davidson 1967a/2002.

us *not* to give ‘kissed’ these surface semantics. Instead, the analyses use the same underlying three-place predicate, ‘_ was a kissing by _ of _’, both for (29) and for (30). This predicate makes the event of kissing explicit. The prepositional phrase ‘in Tompkins Square’ on this account contributes the predicate ‘_ was *in* Tompkins Square’; the prepositional phrase ‘after midnight’ contributes the predicate ‘it was *after* midnight’ and so on. The event approach solves the problem of productivity by introducing events about which indefinitely many things can be said.

Another distinctive feature of the event analysis is that it describes the semantics of action sentences in terms of an existential quantifier phrase — ‘there is an event’ in (34) and (35). This quantifier takes scope over the entire action description, containing the verb and all its modifiers. This quantificational structure correctly predicts that the verb and all its modifiers describe one and the same event, while no constituent in the sentence explicitly identifies or names that event. As we shall see, the event analysis also draws on this added structure to explain what inferences the resulting sentences support, and how they function in connected discourse.

The formal implementation of this event analysis provides a good illustration of some of the contributions of logic to the practice of semantic analysis. To start, suppose we use the formalism of higher-order logic as a meta-language for doing natural language semantics: then we can start to capture the event analysis of (34) and (35) in the formalism by introducing the symbolism $kissed(E,X,Y)$ to indicate that event E was a kissing by X of Y . This mnemonic is useful not only because it is short, but also because it highlights one of the claims of the theory, namely that whenever the verb ‘kissed’ appears in English it expresses a relationship that underlyingly characterizes an event. The notation suggests, in a way that paraphrases like (34) and (35) might not, that the events need to be made explicit to give any systematic semantics for sentences with the verb ‘kissed’. Letting M be Mary, J be John and T be Tompkins Square, we now have these symbolic representations for the truth conditions of (29) and (30):

36. $\exists E (kissed(E,M,J))$.

37. $\exists E (kissed(E,M,J) \wedge in(E,T))$.

By symbolizing the contribution of ‘in Tompkins Square’ using the logical term $in(E,T)$ we are able to concisely indicate that the kissing event is what is being described and that what is being said is that the event was in the park. Again, the notation emphasizes the claimed analysis, namely that ‘in’ always expresses a relationship between a thing (in this case an event) and a spatial landmark that locates it.

We can develop the formalism further to show the event analysis explains why we can keep adding prepositional phrases to modify verbs of change without a loss of grammaticality; it does so by treating such additions through the mechanism of predicate conjunction. Predicate conjunction is a ubiquitous semantic operation (Partee and Rooth, 1983, Chung and Ladusaw 2003); you can clearly see the need for it in examples such as (38)-(40).

38. Felix is a round table.

39. Felix is a brown round table.

40. Felix is a wooden brown round table.

We could continue to add adjectives without a loss of grammaticality indefinitely. It is standard to treat this sort of modification as conjunctive. (38) is true just in case Felix is both round *and* a table. So even though grammar legislates that ‘round’ in (38) is an adjective that modifies ‘table’, since (41) paraphrases (38),

41. Felix is a table *and* Felix is round.

where ‘table’ and ‘round’ overtly describe Felix in (41), they must do so covertly as well in (38). If you understand ‘round’ and ‘table’, all you need to learn to understand ‘round table’ is that it applies to an object just in case both the simpler component expressions apply to it as well. The general rule from which this specific case derives is (PA).

(PA) A complex common noun devised by prefacing a common noun (itself possibly complex) with an adjective applies to an object just in case *both* the adjective *and* the common noun apply to that object.

So (PA) applies to (40) three times, and tells us that overall (40) is true just in case ‘wooden’ and ‘brown’ and ‘round’ and ‘table’ all apply to Felix. The same kind of mechanism applies in composing (36) and (37). We first build a complex predicate characterizing an event — that it was a kissing by Mary of John, that it was in Tompkins Square — and finally derive the sentential condition that there is some event that satisfies this complex predicate.

Using a logical vocabulary also helps us demystify these compositional derivations. A complex conjunctive statement — ‘John is tall and Mary is happy’ — is true if and only if its simpler conjuncts — ‘John is tall’, ‘Mary is happy’ — are true. Letting $tall(J)$ and $happy(M)$ represent the content of the individual conjuncts allows us to represent the content of the conjunctive statement by $tall(J) \wedge happy(M)$. Formalism in our meta-language here highlights our commitment to ordinary logical conjunction. For example, it allows us to carry out inferences on our representation using the deductive operation of conjunction reduction. From the fact that a sentence of the form $\lceil P \wedge Q \rceil$ is true, we can infer that ‘P’ is true and we can infer that ‘Q’ is true.

Analogously, a complex conjunctive predicate — ‘is tall and happy’ — is true of an object if and only if its conjuncts — ‘is tall’, ‘is happy’ — are as well. Now, higher-order logic lets name predicates using the formalism of λ -abstraction: $\lambda X. tall(X)$ names the predicate that is true of an object X if X is tall. In general $\lambda X. C$ names the function which gets its result on argument A by evaluating C on the assumption that $X=A$. This equivalence is embodied in the λ -calculus equivalence $(\lambda X. C)(A) \equiv C[X \leftarrow A]$. For example, $(\lambda X. tall(X))(f) \equiv tall(f)$.

With this notation, we can represent the conjunctive predicate with the λ -term $\lambda X. tall(X) \wedge happy(X)$. In general, we can represent the conjunction of the predicates F and G by the λ -term $\lambda X. F(X) \wedge G(X)$. Thus this notation allows us to rewrite (PA) compactly in a way that shows its close similarity to the interpretive rules illustrated in (5)-(7). We give this new rule as (PAL).

(PAL) If N is a (complex) common noun that refers to the predicate G and A is an adjective that refers to the predicate F , then $\lceil A N \rceil$ is a complex common noun that refers to the predicate conjunction $\lambda X. F(X) \wedge G(X)$.

This formalism again allows us to highlight our commitment to ordinary logical conjunction. Ultimately the same meaning of ‘and’ is at play whether we are conjoining sentences or predicates. For example, the expected inferences follow: If a complex predicate $\lceil \lambda X . P \wedge Q \rceil$ is true of some object A , in other words $\lceil (\lambda X . P \wedge Q)(A) \rceil$ is true, then $\lceil P[X \leftarrow A] \wedge Q[X \leftarrow A] \rceil$ follows (by λ -calculus equivalence) and so both $\lceil P[X \leftarrow A] \rceil$ and $\lceil Q[X \leftarrow A] \rceil$ follow (by conjunction reduction) and thus both $\lceil \lambda X . P \rceil$ is true of A and $\lceil \lambda X . Q \rceil$ is true of A (by λ -calculus equivalence, again).

Now, according to the event analysis, the conjunctive predicate ‘was a kissing by John of Mary *and* was in Tompkins Square’ is true of an event just in case its simpler component predicates ‘was a kissing by John of Mary’ and ‘was in Tompkins Square’ are both true of that event as well. Formally, we represent the conjoined predicate as $\lambda E . kissed(E, M, J) \wedge in(E, T)$ and derive it by predicate conjunction from $\lambda E . kissed(E, M, J)$ and $\lambda E . in(E, T)$.

Another key ingredient of the proposal is the existential quantification over events (in both (34) and (35)) in their main verb and (in (35)) in its preposition. The latter two sentences tell us to treat the former two as existentially quantifying over an event. An initial response to this suggestion might be doubt, since (29) and (30) do not seem to *assert* the existence of an event; i.e., they don’t seem to harbor any existential quantifier. But in fact, the patterns of inference of action sentences are exactly what you would expect from the existential analysis. For example, (42) does not entail (43), but (43) does entail (42):

42. Mary kissed John in Tompkins Square and Mary kissed John on Wednesday.

43. Mary kissed John in Tompkins Square on Wednesday.

This is just what you would expect if the meaning of (42) is represented with two existential quantifiers as in (44) and the meaning of (43) is represented with one as in (45).

44. $\exists E (kissed(E, M, J) \wedge in(E, T)) \wedge \exists E (kissed(E, M, J) \wedge on(E, W))$

45. $\exists E (kissed(E, M, J) \wedge in(E, T) \wedge on(E, W))$.

Patterns of discourse connectivity are also exactly what you would expect from the existential analysis. The familiar pronoun ‘it’ in (46) seems to be about something, presumably an action.

46. John buttered his toast and he did *it* after midnight.

It is reasonable to ask whether the pronoun ‘it’ in (46) anaphoric or deictic? Since (46) is intelligible without a concomitant pointing gesture or any other sort of demonstration, ‘it’ must be anaphoric. But what is ‘it’-s anaphor? The other noun phrases in (46) are ‘John’, ‘his toast’, ‘he’ and ‘midnight’, but these are not actions, and so ‘it’ neither refers back to nor is bound by any of them. Since nothing else is explicitly referenced in (46), we must understand its pronoun ‘it’ to be anaphoric. But what can its anaphor be? The event approach provides a simple answer: it interprets the verb ‘butter’ in (46) as having a place quantifying over events. (46) is understood as (47).

47. Some event was a buttering of the toast by John and *it* (i.e., the event) occurred after midnight.

(47) explains why 'it' in (46) is intelligible. It is an anaphoric pronoun whose anaphor is 'some event'. The event approach posits an *implicit* position in verbs *and* the behavior of 'it' in (46) provides the right evidence for this posited place. The support is unexpected and intuitively pleasing. Data of the sort provided by (46) reminds us that part of what we learn in learning the meaning of a verb of change (indeed, any verb) is its adicity, and what sorts of entities the expressions which fill these places are true of. Verbs of change, it turns out, have an event place.

As explanatorily rich as the event account is it still faces problems. As developed thus far it yields unwanted inferences if applied generally. To see this, consider the distinction between the modifiers 'quickly' and 'allegedly'. The adverb of manner 'quickly' in (48) is modifying the verb 'drove'.

48. Frank drove quickly.

The manner in which the adverb modifies the verb in (48) renders it impossible for (48) to be true and (49) false.

49. Frank drove.

This sort of inference from a sentence with adverbial modification to one without it holds in general. The inference is not only valid but valid *in virtue of its form*, much like inferential relations among (29) through (33).

By contrast, if we symbolize (50) along the lines recommended by the event approach, it winds up logically implying (51).

50. Clinton *allegedly* lied under oath.

51. Clinton lied under oath.

The interpretation of (50), according to the event approach, is (52).

52. There is an event *e* such that *e* was a lying under oath by Clinton *and* *e* was alleged.

But alleging that something is all by itself doesn't make it so. This may be a problem for the event approach, for it focuses on a somewhat special class of modifiers, 'quietly' as opposed to 'allegedly', 'in Tompkins Square' as opposed to 'in a dream'.

Should we ditch the event approach on this basis and seek a more general approach? Before sacrificing it, a closer look at so-called 'problem' modifiers, for example, 'allegedly', suggests a non-*ad hoc* way of distinguishing problem cases from non-problem ones. We can move from (53) to (54) without a change in meaning.

53. Oedipus allegedly married his mother.

54. It has been alleged that Oedipus married his mother.⁸

Since (53) and (54) are synonymous, we can infer that 'allegedly' is not modifying the verb 'married' in (53), but rather, as in (54), 'it has been alleged that' is predicated of a statement. This makes sense, since it is statements that are alleged. And, so this transformation works, because a modifier like 'allegedly' is not an *adverb* modifying a verb, but instead it modifies whole sentences. To take another example, when you say 'John *clearly* studied for this exam' you are not attributing clarity to his extended event

⁸ 'Supposedly John left' ('It is supposed that John left'), 'John killed Bill in Fred's dream' ('Fred dreamt that John killed Fred').

of studying for this exam. Rather, you are asserting ‘*It is clear that* John studied for this exam’. What is clear is the statement that John studied for this exam. In this regard, adverbs like ‘allegedly’ and ‘clearly’ differ from ‘quickly’ and ‘quietly’. Transforming (48) into (55) does not preserve grammaticality.

55. *It has been quick that Frank drove.

If correct, we may have a principled distinction between adverbs the event approach accommodates and those it cannot. What we do not know yet is, if not as conjoined predicates attributed of events, what sort of contribution do adverbs like ‘allegedly’ and ‘clearly’ make to sentences in which they occur. Whatever contribution it is, it will be distinct from the sort of contribution an adverb like ‘quickly’ makes, for suppose that Frank is identical to your uncle, then if Frank drives quickly, so too does your uncle. But even though Oedipus’ mother was Jocasta, no one alleged that Oedipus married Jocasta. If this is right, then adverbs like ‘allegedly’, ‘clearly’, ‘intentionally’, ‘on purpose’ cannot be treated as predicates of events or actions at all. What, then, is their contribution to meaning? We cannot hope to settle all these issues here.⁹

4.2 *Limits to analysis*

It is unchallengeable that the event approach is a paradigm of a substantial analysis about a systematic compositional meaning theory for natural language. Even though these gains are nothing to sneer at, it is an unsettling exchange we have transacted—one that highlights the holistic, scientific character of our investigation into meaning. In effect, we have bargained that if the event approach explains various logical relations and accounts for linguistic productivity, we will ignore our intuitions about the sentences in question. We will ignore our intuition that the seemingly two-place verb ‘_ kissed _’ in (29) through (33) is elliptic for a three-place verb ‘_ was a kissing by _ of _’. Granted, this consequence is not as bad as the discredited proposal that left the adicity of each verb of change as an open-ended affair. But even though on the event analysis there is a fixed number of places for each verb of change, the number that the approach posits is not what we pre-theoretically think it is. Also, whatever contribution prepositions make to the meaning of sentences in which they occur, we do not pre-theoretically think that it’s the same sort of contribution that predicates make. Yet, on this proposal, prepositions in (29) through (33) turn out to be two-place predicates (‘in’ in (30) through (33) is elliptic for the two-place predicate ‘is in’; and so forth). And, lastly, prior to encountering the event approach, it would never have occurred to an English speaker to conclude that (29) through (33) harbor hidden existential quantifiers. So in order to endorse the event approach we had to defend these three surprising and unintuitive consequences. In the end, we can muster a wide variety of evidence in support of these consequences of the event analysis — despite its surprises, the event analysis offers the best account available.

Many proposed analyses in semantics are not such clear wins. For example, analyses that purport to give logical equivalents for lexical meanings are particularly problematic. Paradigmatic examples of such analyses include the hypothesis that (56) is synonymous with or elliptical for (57), that (58) is synonymous with or elliptical for (59), and that (60) is synonymous with or elliptical for (61).

⁹ For further discussion see Parsons 1990.

- 56. Chris used stain to darken the wood.
- 57. Chris used stain to cause the wood to become darker.
- 58. Kim boiled the water.
- 59. Kim caused the water to boil.
- 60. The tiger killed the elephant.
- 61. The tiger caused the elephant to die.

At first, these analyses might seem to offer comparable advantages to the event analysis. For one thing, these analyses immediately account for some important inferential relationships among sentences. It's clear, for example, that (56) definitely entails (57), that (58) definitely entails (59), and that (60) definitely entails (61). For another thing, these analyses, like the event analysis, explain apparently productive processes in language. In English, the '-en' suffix that appears on 'darken' works more generally to turn an adjective *X* into a verb meaning, roughly, 'make more *X*'. It applies to lots of adjectives (though only to short Anglo-Saxon roots ending in consonants of low sonority). Similarly, a wide variety of verbs describing processes things undergo, like 'boil', can also be used in a different sense meaning, roughly, to cause something to undergo the process.

On balance, however, we don't think the preponderance of the evidence favors these analyses. The productivity of these constructions is limited. There is a qualitative difference between the infinite productivity of open-ended modification and the vast but still finite productivity of a lexical derivation that applies to any root word in the language. Speakers have to learn each root form — root forms are arbitrary pairings of sound and meaning, after all — so it's certainly possible that speakers could learn something about each derived form, too. In fact, it's common for complex words to have *motivated* meanings that are related to the meanings of the parts but not predictable from them. The noun 'wheelchair' is a particularly clear case: the word contains the words 'wheel' and 'chair', and, as you might expect, all wheelchairs are chairs on wheels. Yet 'wheelchair' doesn't just mean chair on wheels: it refers to a specific common kind of chair with wheels. It is a reasonable hypothesis that the verb meanings for 'darken' in (56) and 'boil' in (58) are motivated in this sense: they contain a morpheme that expresses causality, but have a more specific meaning that reflects this causality in this semi-compositional way. You can cause something to be darker without darkening it. Turning down the dimmer switch on a light bulb makes the bulb darker but we don't call that 'darkening the bulb' in English. Lowering the pressure in a container of water can make the water boil but isn't boiling the water. And it is also possible to cause something to die without killing it; Romeo's suicide causes Juliet to die, but she kills herself.

These mismatches are reminders that providing analyses in terms of covert logical structure is generally a counterintuitive business. Intuitively, killing is killing and dying is dying — different kinds of things, though certainly related. The same goes for the boiling that you do and the boiling the water does, or the darkening that you do and the dark wood that results — these are different kinds of things, though certainly related. In many cases, such relationships can be captured in lots more transparent ways than by analyses that assume that language contains complex hidden structure.

5 Open problems and further reading

In this chapter, we have suggested that philosophy has a place for truths about meaning, and have explored logic as a tool for developing these truths. Throughout, our discussion we have emphasized both the strong points of this program and its limitations.

On our view, truths about meaning can play an important role in the philosophical description of language. We can use our knowledge of meaning to understand sentences and to communicate with one another. However, truths about meaning may have rather less import for the rest of philosophy — for questions in epistemology or metaphysics, for example — because facts about meaning won't come with *a priori* certainty. Of course, the particular facts we discover these facts may help us clarify our philosophical arguments, and while such possibilities are beyond the scope of this chapter, ongoing debates certainly explore them. A case in point is the contextualist semantics for *know* where uses of the verb attribute knowledge not absolutely but only with respect to prevailing epistemic standards (Lewis 1996).

On our view, logic provides an important collection of tools that we can take as models for the semantic analysis of natural language. The logical treatment of sentential operators, functions and arguments, variables, quantifiers and scope, tense and modality have all inspired research in natural language semantics. There are many good introductions to the overlap between logic and language, particularly Gamut (1991), Kamp and Reyle (1993), and Larson and Segal (1995). There is even a *Handbook of Logic and Language* (van Benthem and ter Meulen, 1997). Technical sophistication comes in particularly handy when the theorist aims to establish a direct correspondence between fine-grained linguistic structure and a compositional theory of truth; so logic is a mainstay of formal linguistic semantics. At the same time, however, we have seen how pragmatic dimensions of language use — context dependence and context change — may motivate extensions to or departures from mathematical logic in semantic analysis. *The Handbook of Pragmatics* (Horn and Ward, 2004) provides an entrée into the diverse phenomena that characterize language in use.

Finally, however, we emphasize that semantic analysis is an empirical discipline that aims to reconcile our intuitions about meaning with our abilities to use language, draw inferences, evaluate arguments and communicate ideas. Philosophy of language provides a framework for this science and logic provides a tool — in many respects, this is where the action finally starts. Fundamental issues remain controversial: what kinds of entities are required for semantic theory? How do these entities get hooked up with our words and utterances? What constraints must a semantic theory for natural language respect? And what is the role of these constraints and our own experience in allowing us to acquire our knowledge as speakers of our native language? We have highlighted the opportunities and the challenges of these questions in semantic analysis here, with the example of action verbs. But they are as open-ended as language itself.

6 References

- Carlos Areces, Patrick Blackburn, and Maarten Marx. Hybrid logic: Characterization, interpolation and complexity. *Journal of Symbolic Logic*, 66(3):977–1010, 2001.
- Chris Barker. The dynamics of vagueness. *Linguistics and Philosophy*, 25(1):1–36, 2002.

- David Beaver. *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford CA, 2001.
- David Beaver. The optimization of discourse anaphora. *Linguistics and Philosophy* 27(1):3-56, 2004.
- Johan F. A. K. van Benthem and G. B. Alice ter Meulen, editors. *The Handbook of Logic and Language*. MIT Press, Cambridge, Mass, and North-Holland, Amsterdam, The Netherlands, 1997.
- Sandra Chung and William Ladusaw. *Restriction and Saturation*. MIT Press, Cambridge Mass, 2003.
- Donald Davidson. The logical form of action sentences. In N. Rescher, ed., *The Logic of Decision and Action*, pages 81-95. University of Pittsburgh Press, Pittsburgh, 1967. Reprinted with discussion in Davidson [2002], pages 105-148.
- Donald Davidson. Truth and meaning. *Synthese* 17:304-323, 1967. Reprinted in Davidson [2001], pages 16-36.
- Donald Davidson. *Inquiries into Truth and Interpretation, Second Edition*. Clarendon Press, Oxford, 2001.
- Donald Davidson. *Essays on Actions and Event, Second Edition*. Clarendon Press, Oxford, 2002.
- Gottlob Frege. *Begriffsschrift, eine der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Halle: L. Nebert, 1879. Translated as *Begriffsschrift, a Formula Language, Modeled upon that of Arithmetic, for Pure Thought*. In *From Frege to Gödel*, edited by Jean van Heijenoort. Cambridge, MA: Harvard University Press, 1967. Also as *Conceptual Notation and Related Articles*. Edited and translated by Terrell W. Bynum. London: Oxford University Press, 1972.
- Melvin Fitting and Richard L. Mendelsohn. *First-order Modal Logic*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1998.
- Jerry A. Fodor. *The Language of Thought*. Crowell, New York, 1975.
- Jerry A. Fodor. *Psychosemantics*. MIT Press, Cambridge Mass, 1987.
- L.T.F. Gamut. *Logic, Language, and Meaning*, Vol. 1: *Introduction to Logic*, and Vol. 2: *Intensional Logic and Logical Grammar*. University of Chicago Press, 1991
- H. P. Grice. Logic and conversation. In P. Cole and J. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41–58. Academic Press, New York, 1975. Reprinted in Grice [1989], pages 22-40.
- H. P. Grice. *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass, 1989.
- Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100, 1990.
- Erhard Hinrichs. Temporal anaphora in discourses of English. *Linguistics and Philosophy*, 9(1):63–82, 1986.
- Laurence R. Horn and Gregory Ward. *Handbook of Pragmatics*, (eds.) Basil Blackwell: Oxford. 2004.
- Hans Kamp, and Reyle, U. *From Discourse to Logic*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1993.
- David Kaplan. Demonstratives. In J. Almog, J. Perry, and H. Wettstein, editors, *Themes from Kaplan*, pages 481-566. Oxford University Press, Oxford, 1989.

- Saul A. Kripke. Naming and necessity. In D. Davidson and G. Harman, editors, *Semantics of Natural Language*, pages 253–355. Reidel, 1972. Reprinted as Kripke [1982].
- Saul A. Kripke. *Naming and Necessity*. Harvard University Press, Cambridge Mass, 1980.
- Alice Kyburg and Michael Morreau. Fitting words: Vague words in context. *Linguistics and Philosophy*, 23(6):577–597, 2000.
- Richard Larson and Gabriel Segal. *Knowledge of Meaning: An Introduction to Semantic Theory*. MIT Press, 1995.
- David Lewis. Scorekeeping in a language game. In R. Bäuerle, U. Egli, and A. von Stechow, editors, *Semantics from Different Points of View*, pages 172–187. Springer Verlag, Berlin, 1979.
- David Lewis. ‘Elusive Knowledge,’ *Australasian Journal of Philosophy*, 74: reprinted in *Papers in Metaphysics and Epistemology*, 1996. pp. 549–67.
- Richard Montague. English as a Formal Language. In Visentini, Bruno et al. (eds.) *Linguaggi nella società e nella tecnica*. Milan: Edizioni di Comunità. 189-224, 1970. Reprinted in Montague (1974), 188-221.
- Richard Montague. The proper treatment of quantification in ordinary English. In Hintikka, K.J.J., Moravcsik, J.M.E., & Suppes, P. (eds.) *Approaches to Natural Language*. Dordrecht: Reidel. 221-242, 1973. Reprinted in Montague (1974), 247-270.
- Richard Montague. *Formal Philosophy. Selected Papers of Richard Montague*. Edited and with an introduction by Richmond H. Thomason. New Haven/London: Yale University Press, 1974.
- Reinhard Muskens. Combining Montague semantics and discourse representation theory. *Linguistics and Philosophy*, 19(2):143–186, 1996.
- Terrence Parsons. *Events in the Semantics of English: A Study of Subatomic Semantics*, MIT Press, Cambridge, 1990.
- Barbara H. Partee. Some structural analogies between tenses and pronouns in English. *Journal of Philosophy*, 70:601–609, 1973.
- Barbara H. Partee. Nominal and temporal anaphora. *Linguistics and Philosophy*, 7(3):243–286, 1984.
- Barbara H. Partee and Mats Rooth. Generalized conjunction and type ambiguity. In R. Bäuerle, C. Schwarze and A. von Stechow (eds.), *Meaning, Use and Interpretation of Language*, Walter de Gruyter, Berlin, 361-383, 1983.
- Paul Pietroski. ‘Logical Form and LF’. in Oxford Handbook to Philosophy of Language. (eds.) E. Lepore and B. Smith, Oxford: Oxford University Press, 2006.
- Hilary Putnam. ‘The Meaning of ‘Meaning’’, 1975, in *Philosophical Papers II: Mind, Language, and Reality*, Cambridge: Cambridge University Press, pp. 215-271.
- W. V. O. Quine. ‘Two Dogmas of Empiricism’, in *From a Logical Point of View*, (Cambridge, Mass.: Harvard University Press) 1953, pp. 20-46.
- W. V. O. Quine.. *Word and Object*. Cambridge, Mass.: The Technology Press of M.I.T; New York, 1960.
- Craige Roberts. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy*, 12(6):683–721, 1989.

- John R. Searle. Indirect speech acts. In P. Cole and J. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 59–82. Academic Press, New York, 1975.
- Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457, 1973.
- Robert Stalnaker. Assertion. In Peter Cole, editor, *Syntax and Semantics 9*, pages 315–332. Academic Press, New York, 1978. Reprinted in Stalnaker [1999], pages 78–95.
- Robert Stalnaker. *Context and Content*. Oxford University Press, Oxford, 1999.
- Alfred Tarski. 1935, ‘Der Wahrheitsbegriff in den formalisierten Sprachen’, *Studia Philosophica*, I, 261–405. Translated into English and published under the title ‘The Concept of Truth in Formalized Languages’ in Tarski [1956].
- Alfred Tarski. 1956, *Logic, Semantics and Metamathematics*, Oxford University Press
- Richmond H. Thomason. Accommodation, meaning and implicature. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editors, *Intentions in Communication*, pages 325–363. MIT Press, Cambridge, Mass, 1990.
- Bonnie Lynn Webber. Tense as discourse anaphor. *Computational Linguistics*, 14(2):61–73, 1988.

Rutgers University Center for Cognitive Science
 152 Frelinghuysen Road
 Piscataway NJ 08854-8020
lepore@rucss.rutgers.edu
Matthew.Stone@rutgers.edu

Supported in part by NSF HLC 0308121.