

**Language in Context**  
*Jason Stanley*  
**Introduction**

A conversation involves acquiring and conveying information about the world, immediately and without much conscious reflection. Our linguistic capacity thereby enables us to engage in complex cooperative activities that require rapid information flow between large numbers of people, such as building bridges and superconductors (though unfortunately also waging war). We share our ability to visually perceive with many other species. In contrast, no other species has our advanced linguistic capacity. Any investigation into what makes humans distinctive as a species begins and ends with language-use. The inquiry into the nature of our linguistic capacities is therefore at the very center of the human sciences.

Perhaps because it is so unique among the species, progress in understanding our capacity for gaining and conveying information about the world linguistically has been slower than progress in understanding (say) our perceptual capacities. We have made many inroads in grammar and phonology. But a complete account of the relation between an utterance and the information about the world it conveys is still probably no closer than a complete account of the biology of the species. A complex phenomenon needs to be broken up into many components to have any hope of ultimately being explained. The phenomenon of language use is no different in this regard. My proposal is a contribution to only a small part of an account of the relation between utterances and the information about the world they convey. For example, there are very difficult questions about what the meanings of word are, and even more difficult questions about the relations of words to the meanings they have. Is the meaning of the word “Mars” the planet Mars? Or is it a rule of use for employing the word “Mars” correctly? If the meaning of the word “Mars” is the planet Mars, how did this relation between a word and a planet come into being? My project in this book only indirectly concerns these difficult questions. I will therefore presuppose without question some standard answers. For instance, because we learn something about the planet Mars from an utterance of “Mars is the nearest planet to Earth”, I will assume that the semantic content of “Mars” is the planet Mars. But no defense of this view of meanings will be given here. Furthermore, I have nothing to say,

in the papers that follow, about the thorny question of the relation between individual words and the things in the world that are their meanings. For example, I will have nothing to say about whether the relation between the words “Mars” and the planet Mars is a causal relation, or even a naturalistically acceptable one. I simply assume, for purposes of my inquiry, that the word “Mars” does refer to the planet Mars. My project begins where questions about standing linguistic meaning leave off.

Once we know the standing linguistic meanings of all the words in a sentence, and the grammar of the sentence, there is still a gap between that knowledge and the information about the world conveyed by an utterance of that sentence. For example, an utterance of the sentence “Jack is tall” may convey the information that Jack is tall for a fourth-grade student. Nothing in the standing linguistic meanings of “Jack” or “tall” seems to determine this information. The papers in this book consist of a defense of one hypothesis about the gap between the standing linguistic meaning of the words in a sentence and the information conveyed by an utterance of it. The defense is not complete, and to even make the sort of hypothesis I do in this work is no doubt premature. But concrete, developed proposals are needed at this stage of inquiry. Even if they turn out to be false, we will at least learn something from their failures.

I have said that my project only indirectly bears on foundational questions such as whether meanings are best conceived of as things in the world (such as the planet Mars) or rules of use. It should be clear that it does not *directly* bear on this question, since I take it for granted that the standing linguistic meanings of non-context sensitive words are often objects, properties, or events in the world. For example, I take it for granted that the standing linguistic meaning of “is taller than six feet” is the property of being taller than six feet, which some people have and others lack. But my project does *indirectly* bear on the question whether we can appropriately explicate meaning in terms of reference to objects and properties in the world. Ordinary language philosophers have argued strenuously that properties such as truth and reference do not apply to linguistic expressions but are rather properties of what people *do* with linguistic expressions. An object is not the meaning of a singular term, and a proposition (an entity capable of truth or falsity) is not the meaning of a sentence. Rather, it is a *use* of a singular term by a person that refers, and it is an *utterance* of a sentence that expresses a proposition; one

cannot speak of a *term* having reference, or a *sentence* having a truth-value, even derivatively. In short: words do not refer, people do. If semantic properties such as referring to an object or being true are not properties of linguistic expressions, then giving an account of linguistic meaning in terms of reference and truth is fundamentally misguided. Philosophers such as Carnap and Tarski were right to focus their attention on formal languages, because the kind of account of meaning they were trying to give (in terms of reference and truth) was inapplicable to natural languages.

Since referring to an object is not property of a linguistic expression, and linguistic meanings are properties of linguistic expressions, ordinary language philosophers sought an alternative account of linguistic meaning. According to it, the linguistic meaning of an expression is a *rule for its proper use*. The motivation ordinary language philosophers gave for the view that linguistic meaning was best explicated in terms of rules of use centrally involved *context-sensitivity*. The fact that one and the same sentence could be used to convey quite different information in different contexts of use seemed incompatible with giving a theory of meaning in terms of reference and truth. First, ordinary language philosophers argued that the existence of context-sensitive words posed a problem for accounts of word meaning in terms of reference. As Peter Strawson writes in “On Referring”:

If someone asks me the meaning of the expression “this”...I do not hand him the object I have just used the expression to refer to, adding at the same time that the meaning of the word changes every time it is used. Nor do I hand him all the objects it ever has been, or might be, used to refer to. I explain and illustrate the conventions governing the use of the expression. This *is* giving the meaning of the expression.

Secondly, the gap between the kinds of accounts of meaning in terms of reference and truth and the intuitive content of utterances seemed large enough as to vitiate the enterprise. For example, for a long time, the only model of the meaning of sentential connectives such as “if...then” and “or” was in terms of truth-functions. Yet the ordinary language meaning (and context-sensitive nature) of these sentential connectives was not adequately captured by truth-tables.

There were two kinds of responses to ordinary language philosophers on behalf of philosophers who thought of meaning as best explained in terms of reference and truth. First, in his extraordinarily influential paper “Logic and Conversation” (Grice, 1989a), the philosopher Paul Grice set out to defend the truth-table analysis of the meaning of the natural language sentential connectives from the ordinary language onslaught. Ordinary language philosophers often did not distinguish the *truth* of an utterance from the *acceptability* of that utterance. The key to Grice’s defense of the truth-table analysis of the meanings of “and”, “or” and “if...not” is that these notions can (and often do) come apart. A given utterance can be true, even though uttering it is not acceptable, because it violates conversational norms. In explaining this distinction, Grice provided the foundations for a theory of conversational norms. The theory Grice gives clearly explains how an utterance may be true, though unacceptable as an assertion due to specific facts about the conversation and its participants. Grice then used the distinction between the truth of a statement and its conversational acceptability in a defense of the thesis that the connectives of propositional logic were correct explications of their natural language counterparts. More specifically, Grice assumed that the natural language logical particles have the truth-table meanings of their logical counterparts, and argued that features of the uses of these expressions that are not explicable by the truth-tables are due to facts about the norms governing conversation, rather than the meanings of the words.

Though most philosophers no longer accept many aspects of Grice’s positive account of the meaning of natural language sentential connectives, everyone accepts that some gaps between what an utterance says and what it conveys are due to general conversational norms, rather than facts about meaning. Thus, appeal to such gaps cannot be used to undermine an account of meaning in terms of reference and truth. The fact that an instance of “P or Q” is only felicitous when the person uttering it does not know that p and does not know that q cannot be used as a premise in an argument for the conclusion that “or” does not refer to a truth-function. For this fact about “or” is due to the general conversational norm that one should convey as much information as one knows. So if one knows that p, one should assert that p, rather than asserting the weaker proposition that either p or q. If Hannah knows that John was at the party, it would be a violation of the Grice’s *maxim of quantity* for her to assert that either John was at the party or he was at

home. She would not be being maximally informative by asserting the disjunctive statement, and hence would be violating the maxim of quantity. The oddity of Hannah's utterance is fully consistent with, and indeed explained by, the fact that "or" refers to a certain truth-function.

The second kind of response to the ordinary language philosopher emerged from the increasing sophistication of semantic theories. Recall that Strawson ridiculed the notion that one could use the notion of reference to give an account of the meaning of a context-sensitive word, such as "this". But the philosopher Richard Montague, and his students Hans Kamp and (in particular) David Kaplan, developed semantic theories where the truth-conditions of a sentence were given *relative to a context of use*. For example, on the semantic theory developed by David Kaplan in his seminal paper "Demonstratives" (Kaplan, 1989), it makes perfect sense to speak of a context-sensitive singular term having reference, albeit relative to a context of use. The context-independent meaning of a context-sensitive singular term, such as the English first-person pronoun "I", is modeled, on such theories, as a *function* from contexts to persons. For example, the context-independent meaning of "I" is a function that takes a context, and yields as a referent the person speaking in that context. So when Bill Clinton utters "I am tired", he expresses the proposition that Bill Clinton is tired, because the context-independent meaning of "I", as applied to a context in which Bill Clinton is speaking, yields Bill Clinton as value. On this approach, it makes perfect sense to attribute reference and truth to expression types, once contextual relativity is factored into the semantic theory. Whereas the notion of a rule of use is vague and mystical, Kaplan's notion of the standing linguistic meaning, or (as he calls it), the *character* of an expression, is not only clear, but set theoretically explicable in terms of fundamental semantic notions; the character of an expression is a function from a context to the reference of that expression in that context. Far from context-sensitivity being an impediment to giving a proper account of linguistic meaning in terms of reference and truth, appeal to these semantic notions allows us to give a considerably more explicit characterization of linguistic meaning than the ordinary language philosophers were capable of providing.

So Grice, on the one hand, and Montague and Kaplan, on the other, together provided a systematic defense of the view that meaning is best explicated in terms of reference, rather than rules of use. If Montague and Kaplan are right, it makes perfect sense to attribute reference to word-types, albeit relative to a context. But these philosophers did not thereby defeat the ordinary language argument for the conclusion that truth and falsity cannot be properties of sentences, but only of acts of uttering those sentences, and the propositions such acts express. For even given an account of the standing meaning of context-sensitive words in terms of functions from contexts to referents, there is still a large gap between the content a sentence appears to have, relative to a context, and the information it intuitively conveys. Furthermore, this gap cannot be explained away by general conversational norms.

For example, consider again the sentence “Jack is tall”. Uses of this sentence can express different propositions, relative to different contexts of use. One use of the sentence can express the proposition that Jack is tall for a fourth-grader. Another use of the sentence can express the proposition that Jack is tall for a basketball player. But general conversational norms surely do not explain the gap between the linguistically determined content of “Jack is tall”, relative to a context of use, and the proposition it intuitively is used to express. Indeed, the linguistically determined content expressed by “Jack is tall”, even relative to a context, does not seem susceptible of truth and falsity at all. It is only when we add the additional contextually provided information –tall relative to *what*, that we are capable of assessing it for truth and falsity. This suggests that the ordinary language philosophers were correct in maintaining that properties such as truth and falsity cannot be said to belong to sentences, even relative to a context. Perhaps, then, meanings are after all better thought of as rules for expressing contents in context, i.e. rules of use.

My project in the papers in this book is to respond to this last remaining vestige of the ordinary language philosopher’s challenge. If my proposal is correct, there is no gap between the linguistically determined content of a sentence, relative to a context, and the proposition it intuitively seems to express. General conversational norms, plus an account of meaning in terms of reference and truth (a *truth-conditional semantics*), explains the

gap between grammar and what is conveyed. I now turn to a more detailed explanation of my project.

Suppose that Hannah utters a sentence in a conversation, say the sentence “Bill Clinton lived in Arkansas.” In uttering this sentence, Hannah imparts certain information about the world. Intuitively, the information that Hannah imparts to her interlocutors is that a certain person, Bill Clinton, lived in Arkansas. As native speakers of the language, we know that this is what has to be true about the world, in order for Hannah’s utterance of the sentence “Bill Clinton lived in Arkansas” to be true. Accordingly, we may call this information the truth-conditions of Hannah’s utterance. Ordinary speakers have fairly robust intuitions about the truth-conditions of an utterance. For example, consider an utterance of the sentence “Aristotle was a great philosopher.” We all have the intuition that what this utterance says would be false, relative to a situation in which Aristotle had died in infancy. The conditions for the truth of what is said by an utterance of “Aristotle was a great philosopher” are simply not met relative to a situation in which Aristotle died in infancy. Similarly, consider an utterance of “John is tired and Sue is tired”. Is what is said by this utterance true relative to a situation in which only John is tired? Clearly, it is not. The truth-conditions of an utterance of “John is tired and Sue is tired” require both John and Sue to be tired. As native speakers of the language, we have robust intuitions about the truth and falsity of what is said by an utterance of English relative to different possible situations.

Speakers have robust intuitions about the grammaticality and ungrammaticality of sentences of their native languages. This fact has been used by linguists to help construct theories of the grammars of natural languages; the grammaticality intuitions we have about our native tongues form the evidential basis for the hypotheses of syntactic theory. The fact that native speakers of a language have robust intuitions about the truth and falsity of what is said by utterances of sentences of that language has been similarly exploited by theorists across a number of disciplines interested in developing theories of meaning for natural languages. In philosophy of language, theorists have exploited these intuitions to argue for many conclusions about meaning. For example, Saul Kripke has exploited such intuitions to argue that names are not disguised definite descriptions, as Bertrand Russell had claimed. In semantics, linguists have exploited these intuitions to

develop accounts of the meaning of a wide range of constructions. In short, intuitions about the truth and falsity of what is said by utterances of sentences have formed the data by which theorists have tested their hypotheses about meaning. There is no other obvious source of native speaker intuitions that are related to meaning. So if we did not have robust intuitions about the truth-conditions of our utterances, it would not be clear how to test such hypotheses; there would be no firm basis on which to construct a theory of meaning.

Native speakers possess intuitions about the truth and falsity of what is said by utterances of sentences of their native languages, relative to various possible situations. This fact has been exploited to construct accounts of the meanings of particular constructions in natural language, and indeed it is the only viable basis for constructing an account of meaning for a natural language. But what is the relation between an utterance and its intuitive truth-conditions? Or, more specifically, how does the utterance come to be associated with those truth-conditions?

It is natural to explain how a specific utterance comes to be associated with the truth-conditions it seems to have (to a rational speaker in good communicative conditions) by adverting to the meanings of the constituent words in the sentence uttered, together with their occurrence in the grammatical structure of that sentence. Perhaps an utterance has the truth-conditions it seems to have because the words in the sentence being uttered have a certain meaning, and they are combined in a certain way in the sentence. When so combined, these meanings determine the truth-conditions an utterance of that sentence seems to have. An utterance of the sentence “Bill Clinton lived in Arkansas” has the intuitive truth-conditions it does because the meanings in English of “Bill Clinton”, “lived”, “in”, and “Arkansas”, when combined in accord with the grammar of the sentence, uniquely determine these truth-conditions. An utterance of the sentence “Bill saw Hannah” has distinct truth-conditions from an utterance of “Hannah saw Bill” because the two sentences, while containing the same words, contain them in distinct grammatical positions.

No doubt, the fact that the words that occur in a sentence have the meaning and grammatical position they do plays a substantial role in any explanation of why an utterance of that sentence has the intuitive truth-conditions it does. But this cannot be the

entire explanation of why utterances have the intuitive truth-conditions they do. For it is uncontroversial that these features of a sentence underdetermine the truth-conditions of an utterance of that sentence. For example, the English sentence “I am human” has a single grammatical structure (i.e. it is not structurally ambiguous) and it is composed of words, each of which has a single, unvarying linguistic meaning. But different utterances of this sentence intuitively have different truth-conditions. If Bill Clinton utters this sentence, the intuitive truth-conditions of his utterance are that Bill Clinton is human, whereas if I utter this sentence, the intuitive truth-conditions of my utterance are that Jason Stanley is human. So it cannot be that the linguistic meanings of the words used in a sentence together with their grammatical positions provide a complete explanation of why utterances have the intuitive truth-conditions they do.

The fact that the words in the sentence being uttered have the linguistic meanings and grammatical position they do therefore is part of the explanation of why that utterance has the intuitive truth-conditions it does, but it is not a full explanation. Facts about the extra-linguistic context in which the sentence is uttered help determine its intuitive truth-conditions in that context. The investigations in this book are devoted to explaining how extra-linguistic context, together with linguistic meaning and grammatical structure, gives rise to content.

A number of theorists are pessimistic about the project of explaining how extra-linguistic context interacts with linguistic meaning and structure to yield the intuitive truth-conditions of an utterance; they hold that beyond linguistic meaning and structure, there is nothing systematic to be said about the relation between an utterance of a sentence and its intuitive truth-conditions. On this view, there are ever so many reasons why utterances have the intuitive truth-conditions they do, and nothing of great import to say about what unifies these disparate explanations. At best, what a theorist can do is say what the linguistic meaning of the words pronounced in a sentence are, and describe the grammatical structure of the sentence. Call this broad range of views, the *pessimistic view*. In contrast, one who adopts an optimistic attitude towards the study of the relations between utterances and their intuitive truth-conditions holds that there *is* an illuminating uniform description of the relation between utterances and their intuitive truth-conditions,

beyond just a specification of context-independent linguistic meaning and grammatical structure.

On the pessimistic view, there is stability to word meaning and the significance of the syntactic structure of sentences. But in general there is no systematic way of going from the meanings of the words in a sentence and its syntactic structure to the intuitive truth conditions of its various utterances. Traditionally, one motivates this position by arguing that the unsystematic way in which the context of discourses affects the intuitive truth-conditions of the utterances that occur in them undermines the search for an illuminating general description of the mapping between utterances and their intuitive truth-conditions. According to Noam Chomsky, for example, there is no scientific study (i.e. no interesting systematic account) of the relation between utterances of sentences and the truth-conditions they intuitively possess. There is still, according to Chomsky, a science of meaning. But such a science cannot investigate the relation between language and the world, since there is no systematic account of this relation (or there is no such relation). Instead, it must involve relations between our sentences and their “I-Meanings”, which, on Chomsky’s view, are mental entities of some kind. Chomsky’s position has been ably defended in recent years by philosophers such as Paul Pietroski.

If any version of the pessimistic view were correct, significant facts about linguistic communication would be inexplicable. From an utterance of a sentence, one gains information about the world. If Hannah utters to Esther the English sentence “There is some chocolate in the kitchen”, and if Esther wants chocolate, she will go to the kitchen. Even if Esther has never heard that particular sentence before, if she speaks English, she will understand what information it conveys about her physical environment. Though the pessimist gives us an account of the relation between Hannah’s utterance and its I-meaning, what spurs us to action is the information the utterance conveys about the world. If the pessimist is correct, there is nothing illuminating and systematic to say about *how* language conveys information about the world. The pessimistic view is difficult to accept because language users (even at very young ages) smoothly grasp information about the world from sentences they have never previously encountered. Furthermore, they do so given only knowledge of the meanings of a (relatively) small number of individual words and their modes of combination into sentences. If there is no systematic

way of proceeding from knowledge of the extra-linguistic context and knowledge of the meaning of individual words and modes of combination into sentences to a grasp of the information about the world that is conveyed by an utterance of a sentence, it is mysterious how language users could so smoothly move from linguistic comprehension to action. In short, if the pessimistic view were correct, the connection between speech and action would be inexplicable.

Furthermore, if the pessimistic view were correct, it would not be clear how to construct a theory of meaning. Native speakers have intuitions about the truth and falsity of what is said by an utterance, relative to various possible situations. But native speakers do not have intuitions about meaning. Even dictionaries are a poor source for discovering the “I-meanings” postulated by theorists such as Chomsky. In short, theorists of meaning who do not think there are systematic relations between the intuitions native speakers have about the truth and falsity of utterances and the meanings of words and sentences have stripped themselves of any plausible evidential basis for their hypotheses.

If the pessimistic view is incorrect, then there must be some way to account, in a systematic manner, for the apparently significantly diverse ways in which linguistic interpretation is conditioned by context. That is, it is incumbent upon those who reject the pessimistic approach to produce a systematic account of the data to which advocates of pessimism have drawn our attention over the years. Meeting this challenge is complicated by the different strategies available for pursuing a systematic account of linguistic interpretation.

One possible way to pursue a systematic strategy for explaining the relation between utterances and the truth-conditions they intuitively possess is to attempt a systematic explanation of rational communicative action generally. On this model, advocated for example by *relevance theorists* such as Dan Sperber and Deirdre Wilson (1986) and Robyn Carston (2002), there is a systematic account of interpretation of the mental states of others. It is our ability to make reasonable inferences from the observable behavior and characteristics of our interlocutors to their mental states that ultimately explains our success in acquiring information from their utterances. On this model, linguistic behavior is no different in kind from the other cues (winks and nods, half-smiles) our interlocutors give about their inner narratives. The way we draw out

information from others is by interpreting, in the first instance, their mental states; their utterances are just evidence, no different in kind from the evidence provided by the kind of clothing they tend to wear.

Advocates of this strategy do think that the meanings of the words in a sentence help us in deciding what the speaker intends to communicate by her utterance of that sentence. But they also think that the linguistically guided information provided by the sentence and its grammatical structure is quite minimal. In interpreting others, we employ non-linguistically guided reasoning about their mental states. This allows us to enrich the linguistically guided information we glean from the sentence they use with other information, via a process advocates of relevance theoretic pragmatics call *free enrichment*, and advocates of alternative pragmatic approaches, such as Kent Bach (1994), call *implicature*. The intuitions speakers have about the truth or falsity of what is said by an utterance is thoroughly influenced by these non-linguistically guided factors.

It is not clear that this strategy can meet the pessimists' challenge. On this view, there is no direct systematic relationship between utterances and the truth-conditions they seem to have. Instead, the systematic relationship is between two interlocutors interpreting one another's mental states. Indeed, there may be no way, on this view, of setting up any kind of systematic relation between utterances and the truth-conditions they seem to have. More worryingly, such a view does not seem to account for the smoothness and rapidity of linguistic communication. When someone hears a sentence containing context-sensitive terms such as "that", "he", or "she" under good communicative conditions, they rapidly grasp the specific truth-conditions of the utterance. Of course, to decode the contribution of the context-sensitive terms, a hearer must appeal to facts about the discourse context, including salient facts about the mental states of her interlocutors. But the fact that we do so quite rapidly suggests that our search for the facts to interpretation is significantly constrained by linguistic meaning, that is, is under linguistic control.

Many philosophers of language, impressed by the challenge of reflecting the pervasive dependence of utterance content on context, have argued that a theory of meaning for a language should not incorporate much context-dependence. According to this *minimalist* approach, the semantic content of a sentence relative to a context bears

only a loose relation to the intuitive truth-conditions of an utterance of that sentence in that context; our intuitions about the latter are the result of both the semantic content and myriad unsystematic pragmatic effects (e.g. Borg (2004), Cappelen and Lepore (2004)). For example, suppose a teacher utters the sentence, “Everyone has passed the exam” to her class. Minimal semanticists maintain that the semantic content of her sentence, in that context, is the proposition that everyone in the universe passed the one and only one exam in the universe. The proposition we intuitively feel is expressed by her utterance is the result of (say) recognizing that this is an absurd proposition to assert, and pragmatically adjusting to arrive at the proposition that everyone in the class passed the exam given in that class. The minimal semanticist postulates informationally impoverished semantic contents of sentences, to avoid admitting context-sensitivity in the theory of meaning.

The philosophers who espouse minimal semantics hope to save systematic semantics from the challenges posed by pervasive context-sensitivity. However, speaker intuitions about the truth and falsity of sentences relative to contexts are the evidential basis for hypotheses about meaning. Minimal semantics severs the data of the theory of meaning from its hypotheses, rendering the semantic project a tapestry of idle speculation. Perhaps for this reason, the project of minimal semantics has few serious adherents in linguistics. Minimal semanticists have had little or nothing to say about how the theory of meaning should be carried out in the absence of any data for its claims.

The popularity of the minimal semanticist’s perspective in philosophy of language is due to the fact that philosophers tend to detect some threat from context-sensitivity to the project of the theory of meaning. However, the nature of this putative threat is far from clear. It is true that some of the progenitors of the analytic tradition in philosophy were suspicious of the use of natural language in reasoning, because of context-sensitivity and ambiguity. But this is because their chief concern was not the theory of meaning, but rather the project of justifying arithmetic. To justify arithmetic, they sought to derive the arithmetical axioms from logical axioms within a system that was plausibly a formalization of pure logic. Carrying out mathematical proofs in a context-sensitive language adds additional risk of fallacy. If a sentence S is context-sensitive, then it can express different propositions, relative to different contexts of use. If the context-

sensitive sentence S occurs multiple times in a proof, say once as a premise and another time as an antecedent of a conditional, then it may express different propositions in these different occurrences. If so, then while the syntactic form of the inference may be valid (in this case, an instance of modus ponens), the inference itself will not be valid. So a context-sensitive language is not an ideal language in which to carry out mathematical proofs; such a language introduces novel possibilities of error. It is for this reason that the progenitors of analytic philosophy, in their attempt to devise an ideal formalism for conducting mathematical proofs, eschewed context-sensitivity. I suspect that much of the aversion philosophers have to context-sensitivity in a theory of meaning is a remnant of a legitimate aversion to context-sensitivity by our philosophical ancestors, given the aims of their projects. But natural language semantics is a very different project than logicism; it is an error to let the concerns of the latter affect us when engaged in the former.

There are additional disadvantages to employing a language rife with context-sensitive constructions. For example, as Timothy Williamson (2005, p. 100) notes, one cost of pervasive context-sensitivity is that it complicates “the preservation of information in memory and its transmission by testimony” (see also Hawthorne (2004, pp. 109-110)). If adjectives such as “tall” introduce a kind of context-sensitivity, say to a comparison class or a degree of height, then in order to transmit the information I convey by an utterance of “Jack is tall”, you must identify the exact comparison class or degree of height I have in mind. Similarly, in order for me to remember what I believed when I believed that Jack is tall I must know exactly what comparison class I had in mind.

So, context-sensitivity complicates preservation of information. It makes communication somewhat more difficult in general, if we think of successful communication as conveyance from speaker to hearer of the very same proposition. If it is a conversational maxim, as some have suggested, that “speakers ought, in general, to assume that their addressees have whatever information is necessary to determine what they are saying” (Stalnaker (1987, p. 110)), then pervasive context-sensitivity will also make adherence to conversational norms more difficult. Would it be better, then, to reject the appearance of context-sensitivity in natural language, and construct a theory of meaning that ignores it? Clearly it would not be. As we have seen, arguments presented in a context-sensitive language face an additional possibility of error. In this sense, giving

arguments in a context-sensitive language is risky. But if context-sensitivity is already present in the language, not recognizing it is considerably riskier. Arguments may seem valid, which in fact are not. Similarly, if we do convey information with the use of pervasively context-sensitive constructions, disregarding context-sensitivity is not a way to preserve the transmission of knowledge by testimony. What would thereby be preserved would not be the information that was in fact originally conveyed by the original utterance, but some irrelevant informationally impoverished monstrosity; contents that are sufficiently impoverished to be faithfully transmitted are not sufficiently interesting to assert. Minimalist semantics is an attempt to retain a fanciful philosopher's illusion, at the cost of exposing us to inferential risk and emphasizing informational absurdity.

Finally, minimal semantics helps not a whit in explaining the problems with which we started. Our problem was to explain how humans rapidly go from sentences they have never before encountered to information about the world. The minimal semanticist can explain how humans rapidly go from sentences they have never before encountered to minimal content. For example, the minimal semanticist can explain how humans rapidly go from hearing the sentence "Everyone passed the exam" to the minimal content that everyone in the universe passed the only exam in the universe. But that is where her explanation ceases. And there are two problems with this project. First, it's not clear that speakers ever do rapidly go from hearing an utterance of the sentence "Everyone passed the exam" to the minimal content that everyone in the universe passed the only exam in the universe; indeed, it's not clear that speakers ever entertain minimal contents at all. The minimal semanticist is therefore providing an explanation for phenomena that never in fact occur. Secondly, even if the minimal semanticist could make it plausible that language users do tacitly entertain minimal contents, they would still be addressing a distinct problem from the one with which we started, which was to explain how language users rapidly go from sentences they have never before encountered to information about the world that spurs them to action. Minimal contents are not the kinds of things that spur language users to action; grasping the proposition that everyone in the universe passed the only exam in the universe is not particularly useful. To provide the missing explanation, the minimal semanticist would have to spend

the bulk of her time developing a sweeping pragmatic account of the relation between sentences and their intuitive truth-conditions, of the sort that relevance theorists or other pragmatic theorists are trying to develop. In other words, even if language users did tacitly entertain minimal contents, minimal semantics would still end up being nothing but a branch of, for example, relevance theoretic pragmatics.

So how do we communicate, if most of our sentences have multiple context-sensitive elements? It may be, as Richard Heck (2002) has argued, that we have been mistakenly seduced into accepting “the naïve conception of communication”, according to which successful communication requires grasp of the same proposition expressed by an utterance of a sentence. Perhaps, as Heck and others have suggested (e.g. Bezuidenhout (1997)), successful communication requires only grasp of a *sufficiently similar* proposition to the one expressed by the utterance. Cappelen and Lepore (2006) rightly protest that it is particularly difficult to give a reductive characterization of the conditions under which two propositions are sufficiently similar in a given context. But perhaps we do not need such a characterization to reject the naïve conception.

I am in fact attracted to a quite different response to the problem than the one entertained by Heck, Bezuidenhout, and others. I am skeptical about the grounds philosophers have given to think that it is very difficult to grasp the proposition expressed by someone else’s utterance of a context-sensitive sentence. First, it is tempting to think that successful communication requires the interlocutor to think of the objects and properties that constitute the proposition in the same way as the person who utters the sentence thinks of those objects and properties. Secondly, it is tempting to think that grasping a property requires being able to distinguish that property from very similar ones. Third, it is tempting to think that one grasps the proposition expressed by someone else’s utterance only if, had one uttered the same sentence in a similar context, one would thereby express the same proposition. But all of these temptations should be resisted.

Suppose someone utters a sentence containing a quantified noun phrase. They thereby express a proposition containing a property that is the domain of that quantified noun phrase. For an interlocutor to grasp that proposition, it is not necessary to think of that property in the same way as the way the person who utters that proposition. For example, suppose I utter the sentence, pointing at John, “He is tired”. I may think of John

somewhat differently than you think of John. But I can still grasp the singular proposition about John that you express. A similar point holds for properties. It is no doubt a difficult matter to state the conditions under which someone has a *de re* thought about an object or a property, and the conditions for *de re* acquaintance with an object or property might themselves be situation dependent. But it is a difficult matter that has nothing specifically to do with context-sensitivity in the philosophy of language.

The availability of this response depends, of course on the conception of proposition that one employs. In giving this response, I have characterized the contents of utterances as Russellian propositions, which are structured meanings that contain objects and properties. If one instead employs a semantic theory with *Fregean* propositions, as Heck advocates, then this response is more difficult to give. Fregean propositions contain ways of thinking of objects and properties. To grasp a proposition containing ways of thinking of objects and properties, it seems one must grasp the specific ways of thinking that constitute the proposition. Explaining how interlocutors can “share content” (to use the apt expression of Cappelen and Lepore (2006)) is, as one might have suspected, more difficult to solve if meanings are individuated in terms of ways of thinking, rather than in terms of objects and properties.

The second thought that leads philosophers to think that grasping the content expressed by someone else’s utterance of a context-sensitive sentence is difficult is that an interlocutor is often unable to distinguish the property intended by the speaker as (for example) the intended domain of quantification from other similar properties. For example, suppose John utters the sentence “Everyone is a philosopher”, meaning to express the proposition that everyone in the room is a philosopher. But suppose everyone in the room is sitting on the right side of the room. One might think that unless John’s audience knows that he intends the property of being in the room *rather than* the property of being on the right side of the room, they do not grasp the proposition he expresses. In short, the thought is that a person does not grasp a proposition expressed by a sentence unless they can distinguish the properties in that proposition from all properties that are sufficiently similar, if they were presented with all the possible alternatives at the same time.

But it is unclear, to say the least, why we should accept the epistemic requirement on grasp of the constituents of a proposition that motivates this line of thought. If someone utters the sentence “Water is wet”, surely I can grasp the proposition thereby expressed, even if I cannot distinguish water from very similar substances. Grasping a proposition surely does not require meeting such a demanding epistemic requirement. If it did, we would have trouble explaining how we grasp the propositions expressed by many eternal (non-context-sensitive) sentences.

The third tempting thought is that one grasps the proposition expressed by someone else’s utterance of a sentence S only if, were one to utter the same sentence in the same or similar context (obviously the speaker of the context would be different), one would express the same proposition. For example, Cappelen and Lepore (2006, section 3.2) write:

Consider two sailors on the ship, Popeye and Bluto. After the sad departure, Popeye observes ‘That was a nice occasion. Every sailor waved to every sailor’. Immediately afterwards, Bluto concurs, ‘That’s right. Every sailor waved to every sailor’. In such a circumstance the following is often obvious: we treat these utterances as expressing agreement. Popeye and Bluto agree that every sailor waved to every sailor...But if Stanley and Szabo [whose theory of quantifier domain restriction is under dispute] were right about the semantics of quantifiers, their concurrence would be a minor miracle.

There are several thoughts contained in this passage. The first is that Bluto grasped what Popeye said only if Bluto’s utterance of the same sentence expressed the same proposition as Popeye’s previous utterance. The second is that the semantics of ascriptions of “agree” requires sameness of content. I will not challenge the second of these thoughts here. But the first seems false.

Suppose speaker intentions determine the value of quantifier domain indices, and suppose, for the sake of argument, that Bluto intends to refer to slightly different domains for his two uses of “every sailor” than Popeye did. Then their two utterances express different propositions. But it simply does not follow that Bluto did not grasp the

proposition expressed by Popeye. Suppose Popeye expressed the proposition that every sailor on the ship waved to every sailor on the shore. So, the domain for Popeye's first use of "every sailor" is the property of being on the ship, and the domain for Popeye's second use of "every sailor" is the property of being on the shore. Bluto could grasp this proposition in this context by thinking of the property of being on the ship (which is the domain for the first quantified noun phrase used by Popeye) as the unique ship-related property intended by Popeye, and he could grasp the property of being on the shore as the unique shore-related property intended by Popeye. In this context, this would suffice to give him a *de re* grasp of the quantifier domain properties contained in the proposition expressed by Popeye's utterance of "every sailor waved to every sailor". So, whether or not Bluto succeeds in expressing the same proposition by his subsequent utterance of the same sentence, he can grasp what is said by Popeye.<sup>1</sup>

So I do not see the case for thinking that ubiquitous semantic context-dependence threatens either the possibility or the systematicity of interpretation. I therefore have sought to develop a systematic account of the apparently unsystematic ways in which the context of discourse may affect the truth-conditions of the utterances that constitute it, one that does not proceed via an appeal to a systematic general theory of rational communicative action.

The account I have been developing and defending is only one possible systematic account of the relation between utterances and the truth-conditions they seem to have. But it is perhaps the most straightforward such account. Consider a sentence *S* as uttered in a context *c*. This utterance of *S* has the truth-conditions it seems to have because these truth-conditions are what results from taking the words contained in *S*, considering what their content is relative to the context *c*, and combining these contents in accord with the syntactic structure of *S*.

---

<sup>1</sup> These thoughts can quite obviously be brought to bear on the conclusions that have been drawn in the literature on indicative conditionals from the example of "Sly Pete", introduced in Gibbard (1981). Sly Pete examples and variants thereof have been taken to undermine the thesis that indicative conditionals have truth-conditions, on the grounds that if they did have truth-conditions, their truth-conditions would vary too much as a function of context. I think similar considerations would show that many more kinds of sentences in our language fail to have truth-conditions, including a wide variety of sentences containing modal expressions.

So, for example, consider the sentence “I am human”, uttered by Hannah. Intuitively, Hannah’s utterance is true if and only if Hannah is human. These are the truth-conditions of Hannah’s utterance, on my view, because they come from taking the *content* of “I” relative to the context in question (which is Hannah), and the content of the predicate “am human” relative to that context (which is the property of being human). Combining these in accord with predicational structure yields the truth-condition that Hannah is human. Similarly, consider the sentence, “She is tired now”, uttered by Hannah at 4:00 p.m. on July 12, 2003 (and suppose Hannah is pointing at Sara). Intuitively, Hannah’s utterance is true if and only if Sara is tired at 4 p.m. on July 12, 2003. These are the truth-conditions of Hannah’s utterance, on my view, because the content of Hannah’s use of “now” is the time of utterance (4 p.m. on July 12, 2003), the content of Hannah’s use of “She” is Sara, and combining this with the content of “is tired” yields those truth-conditions.

We have the ability to gain information about the world from the utterances of others, given only knowledge of the meaning of words and the significance of combining words into more complex expressions. If the explanation of the relation between utterances and their intuitive truth-conditions I have just provided is correct, this ability is not mysterious. It is explained by our grasp of the *contents* (rather than just the linguistic meanings) of individual words and the significance of combining them, for, if I am right, together this yields the intuitive truth-conditions of an utterance.

Of course, some words make different contributions in different contexts to the intuitive truth-conditions of sentences containing them. In fact, the sentences just discussed contained three examples of such words; “I”, “she”, and “now.” Part of learning English involves mastering the context-independent linguistic meaning of these expressions. But mastering the context-independent meanings of these expressions is not sufficient for knowing what they contribute to the truth-conditions of a sentence containing them in a particular context of use. For example, the context-independent meaning of “she” is something like *the indicated (or salient) female*. But in one context, what a use of “she” may contribute to the intuitive truth-conditions of a sentence such as “She is tired now” is one person, whereas in another context, where another person is being indicated, it may contribute another person. Some knowledge of the extra-linguistic

context is needed in order to know the contribution such an expression makes in that context to the intuitive truth-conditions of the sentence containing it. But knowledge of the context-independent meaning of such a term makes the search for such knowledge considerably easier. For example, if a person uses the term “now”, our knowledge of its context-independent meaning restricts our hypotheses about its value to intervals of time containing the moment of that person’s utterance.

My purpose in developing an account of the relation between utterances and their intuitive truth-conditions is to account for two facts. First, we smoothly grasp information about the world from sentences they have never before encountered. Secondly, one and the same sentence may be used to convey different pieces of information about the world, relative to different contexts of use. Both of these facts are unsurprising, given that language is primarily a tool for rapidly conveying information about the world. If we employed a form of communication that lacked context-sensitive elements, then each sentence could only ever be used to express the same piece of information about the world. Such a language, while perhaps helpful for carrying out mathematical investigations, would be impractical for everyday use. It is far easier to utter “that is dangerous”, pointing at a snake in our path, then coming up with a context-insensitive sentence that conveys similar information (perhaps “There is a dangerous snake in front of Jason Stanley and Jeff King”). So for language to be useful, some work has to be done by context-sensitive elements. But if the bulk of intuitive truth-conditions were determined by extra-linguistic context, it would become unclear why we would need to utter sentences at all. My proposal about the relation between utterances and their intuitive truth conditions is a plausible starting hypothesis for how language is able to be sufficiently elastic as to be usable, and sufficiently rule-governed as to be useful.

There are a number of commitments of the particular metaphysical account I defend of the relation between utterances and their intuitive truth-conditions. It is worth being upfront about them here, rather than forcing the reader to discover them over time in reading through the chapters of this volume.

If every term was context-sensitive in numerous different ways, then, consistently with my proposal, interpretation might still not be systematic. One way this could turn out is as follows. If every term in the language had a very thin context-independent

meaning, knowledge of the context-independent meaning would not significantly constrain its interpretation in a context. So, for instance, if the context-independent meaning of “water” was simply *the indicated (or salient) substance*, and the context-independent meaning of “drinks” was simply *the indicated (or salient) relation*, then one would need to know a great deal about the extra-linguistic context in order to interpret a sentence such as “John drinks water regularly.” For it would be extra-linguistic context that one would need to appeal to discover which substance of all possible substances counted as the content of “water” in that context, and which of all possible relations counted as “drinking” in that context. If so, it might turn out that to grasp the intuitive truth-conditions of an utterance, one had to know a great deal about the extra-linguistic context, far more than is plausible, given the smooth way in which we grasp the truth-conditions of sentences we have not encountered before. One commitment of my project is therefore that, for most words in the language, their meaning is considerably more informative than the above meaning-rules for “water” and “drinks.” This is not a terribly controversial claim, at least not in recent philosophy of language.

The second commitment of my project concerns the effects of extra-linguistic context on the contribution words make, relative to a context, to the truth-conditions of sentences containing them. Even if the context-independent meaning of a word like “water” was fairly rich, and excluded substances such as orange juice and petrol from being water, there is still a worry for the systematic nature of interpretation. If context could affect the interpretation of words in such a manner that the content they express relative to a context could be inconsistent with their context-independent meaning, that would threaten the systematic nature of interpretation. For example, suppose that the context-independent meaning of “water” is the substance pure H<sub>2</sub>O. But suppose that one effect of extra-linguistic context was to allow arbitrarily large expansions of the interpretation of “water”, so that quantities of liquid that are not pure H<sub>2</sub>O could count as water. Then, as far as knowledge of meaning goes, an utterance of “This is water” could be true, relative to some context, just in case the indicated quantity is a quantity of any liquid whatsoever. If context could have this effect on interpretation, then the systematic nature of interpretation would be threatened. As far as meaning goes, any sentence could be associated with virtually any truth-condition in some context.

So another commitment I adopt is that extra-linguistic context is never called upon to *expand* the content determined by the context-independent meaning of a term in a context. For example, if a noun N expresses a certain property P, then it is never the case that N together with the addition made by extra-linguistic context contributes, to the truth-conditions of a sentence containing it, a property that is true of more things than P. So, it is plausible to take “tiger” as expressing the property of being a tiger, a certain kind of mammal. A consequence of this second commitment is that there is no context in which the word “tiger”, together with the contributions of extra-linguistic context, contributes, to the intuitive truth-conditions of a sentence containing it, a property such as being a feline, which is true of more things than the property of being a tiger.

The third commitment I adopt involves the significance of syntactic combination. Consider an atomic sentence, such as “Hannah is human.” This sentence predicates “is human” of the content of the name “Hannah.” The semantic significance of predication here may be treated in different theoretical ways. For example, it may be treated as satisfaction (in Tarski’s sense) or functional application (in Montague’s sense) or as property application. So, the truth-condition of “Hannah is human”, relative to a context, is that Hannah satisfies the predicate “is human” or that the function expressed by “is human” maps Hannah onto the true or that Hannah has the property of being human. These different treatments are different ways of explicating the intuitive contribution of predication to the truth-conditions of subject-predicate sentences. But I assume that (in the case of predication) there is one and only one interpretation that is the interpretation of predication in every context. In particular, I assume that the contribution of a syntactic configuration to the intuitive truth-conditions of a sentence containing such a configuration cannot vary as a function of extra-linguistic context, in the way that the contribution of a term such as “she” to the truth-conditions of sentences containing it can vary from context to context.

This last proviso does not exclude the possibility that one syntactic configuration can be ambiguous between two different interpretations, like “bank” is ambiguous between *riverbank* and *financial institution*. What I want to exclude is the possibility that a given syntactic configuration –say the relation between an adjective and a noun in a configuration like “tall woman” or “flat plain”— can have arbitrarily different meanings

in different contexts, in the way that a context-sensitive term such as “she” or “I” can have arbitrarily different meanings in different contexts, consistently with its context-independent meaning. In somewhat theoretical terms, this second commitment rules out the possibility of a syntactic configuration having an *indexical* character.

So, this is my proposal for how an utterance comes to be associated with the truth-conditions it intuitively possesses. Each term in a sentence being uttered has a content that is determined by its context-independent meaning together with extra-linguistic context. The context-independent meanings of most terms in the language are fairly rich, rather than fairly vacuous. The function of extra-linguistic context cannot be to expand the content of a term relative to that context. Finally, these contents are put together into truth-conditions by composition rules determined by the syntactic configuration of the sentence, which are not sensitive to context (except insofar as they may be ambiguous between two or three different composition rules). The result is the intuitive truth-condition of the sentence relative to that context (or the intuitive truth-condition of that utterance).

What results from this picture is a very specific metaphysical picture of the way in which utterances become associated with their intuitive truth-conditions. It is trivial to conceive of possible counter-examples. There are many sentences utterances of which seem to express truth-conditions that cannot be captured in this manner. Let me give some examples.

First of all, consider tense, as in the present-tensed sentence “Bill Clinton is eating lunch.” Intuitively, the truth-conditions of an utterance of this sentence concern the time of utterance. If the utterance of the sentence occurs at 4 p.m. on July 12, 2003, the intuitive truth-conditions of the utterance is that Bill Clinton is eating lunch at 4 p.m. on July 12, 2003. But unlike the sentence “Bill Clinton is eating lunch now”, there doesn’t appear to be a term in the sentence “Bill Clinton is eating lunch”, the value of which, relative to a context, is the time of utterance. If there is no such term, then utterances of present-tensed sentences show that there are some utterances the intuitive truth-conditions of which are determined by more than just the combination of the contents of the terms in the sentence being uttered.

A second example, as we have seen above, involves predicative uses of gradable adjectives, such as “John is tall” or “Hagia Sophia is old.” Intuitively, an utterance of the sentence “John is tall” may express the content that John is tall for a human being, and an utterance of the sentence “Hagia Sophia is old” may express the content that the venerable Byzantine Church is old for a building. But there does not appear to be a term in either sentence the value of which, relative to a context, is the property of being a human, or (in the second case) the property of being a building. So, this is another example of a kind of sentence, utterances of which appear to have truth-conditions that are not determined just by putting together the contents of the words contained in the sentence in accord with their syntactic configuration.

A third example involves the phenomenon of quantifier domain restriction. Suppose John, having just come back from shopping, utters the sentence “Every bottle is in the fridge.” It seems that the result of adding together the content of each word, relative to this context, in accord with the syntactic structure of the sentence, yields the truth-condition that every bottle *simpliciter* is in the fridge. But of course, what John intuitively said was that every bottle *that he just bought* is in the fridge. But the restriction to the things John just bought does not seem to come from the contents of any part of the sentence uttered. So, this is another case where my proposal as to the relation between a sentence and the intuitive truth-conditions of an utterance of it seems incorrect.

A fourth example involves mass-terms, such as “water” or “gold.” It is plausible to take a term like “water” to express the property of being pure H<sub>2</sub>O. If so, then the contribution of “water” to the intuitive truth-conditions of utterances of sentences containing it must be the property of being pure H<sub>2</sub>O. But consider an utterance of “Lake Burley Griffin is filled with water.” Intuitively, this utterance is true. So, its intuitive truth-conditions must in fact be satisfied. But given the above hypothesis about the content of “water”, the truth-conditions of this utterance are not actually met. For it is not true that Lake Burley Griffin is filled with pure H<sub>2</sub>O. The substance that fills Lake Burley Griffin is not pure H<sub>2</sub>O, but has H<sub>2</sub>O mixed with many impurities. So this is another example that appears to threaten a part of my project. For here it seems that the contribution of “water” to the intuitive truth-conditions of utterances of sentences containing it is an ‘expansion’ of its stand-alone content. That is, the contribution of

“water” to the intuitive truth-conditions of utterances of sentences containing it is a property that is true not just of quantities of pure H<sub>2</sub>O, but also of H<sub>2</sub>O mixed with impurities.

There are numerous other examples that suggest that my proposal about the relation between an utterance and its intuitive truth-conditions is overly restrictive. Some may think of this as a reason not to pursue it at all. But my own inclination is to seek an account of the relation between utterances and their intuitive truth-conditions that makes the connection between them explicable. I have great difficulty seeing alternative proposals as sufficiently constrained. Proposals that appear smoothly to account for the bewildering variety of ways in which context appears to affect the relation between utterances and their intuitive truth-conditions only do so because they appeal (whether overtly or covertly) to interpretive processes that are deeply mysterious or wildly unconstrained. I prefer to explore a view that is neither mysterious nor unconstrained, and see how far one can use it to explain our linguistic behavior. If at the end of the day, my proposal is too constrained, perhaps the investigation into its adequacy will reveal systematic ways of modifying or liberalizing it. This seems more promising than beginning with a mysterious and unconstrained process, and trying to add on stipulative constraints. If we don't have a clear grip on the process with which we began, why think that adding on restrictions will yield any greater elucidation?

Fortunately, the alternatives are not “explicable but subject to counterexample” versus “inexplicable, but not subject to counterexample.” The proposal I have sketched above concerning the relation between utterances and their intuitive truth-conditions may seem at first blush to face insurmountable objections. But I argue in this book that a host of prima facie counterexamples to the proposal are not successful. For example, I argue, in each of the above cases, that the intuitive truth-conditions of the utterance are due to the assignment of values to the parts of the sentence uttered, and combination in accord with syntactic structure. In each case, the belief that this cannot be the case is due either to an impoverished conception of syntactic structure, or an impoverished conception of available semantic resources.

Many philosophers have been driven to reject the kind of intimate relation I favor between the result of combining the contents of parts of a sentence in a context and the

intuitive truth-conditions of its utterance because of their antecedent commitments about certain controversial constructions. For example, consider true identity sentences flanked by names, e.g. “Hesperus is Phosphorus” or “Cicero is Tully.” Some philosophers of language hold that the intuitive truth-conditions of propositional attitude ascriptions containing such sentences (such as “John believes that Cicero is Tully”, or “Hannah doubts that Hesperus is Phosphorus”) diverge quite radically from the result of taking the contents of their parts and combining them in accord with their syntactic structure. For such philosophers, the result of the latter procedure, applied to a sentence like “John believes that Cicero is Tully”, yields the same content as the result of the procedure, applied to the sentence “John believes that Cicero is Cicero.” But intuitively, it certainly seems that an utterance of “John believes that Cicero is Tully” could be false, even though John believes that Cicero is Cicero.

I do not, in the papers in this volume, address propositional attitude constructions, and in general (except for the case of quantifier domain restriction and indicative conditionals) I steer clear of constructions that have engendered the greatest philosophical controversy. But that does not mean the proposal is not relevant for these topics. If the proposal works outside the realm of the most philosophically contentious cases, then that provides strong evidence that it is correct in the philosophically contentious cases. Thus the success of this project would provide motivation for investigation views of propositional attitude constructions of the sort defended in Ludlow (1995), according to which there are unpronounced variables in the scope of propositional attitude constructions whose values are modes of presentation, ways of thinking, or other non-extensional entities.

The first paper in this collection, “Context and Logical Form”, introduces my proposal, and argues, in a preliminary fashion, that a wide variety of constructions can be accommodated within its scope. The first section is devoted to explaining the proposal and its commitments. In the second section, I consider the apparent phenomenon of “non-sentential assertion.” Those who reject my account of the relation between utterances and their intuitive truth-conditions often appeal to the fact that we make assertions (thereby presumably expressing truth-conditions) with the use of apparently sub-sentential items. For example, a thirsty man might utter “water” to a street vender, or I might utter “two

apples” to someone selling apples. If such utterances are genuinely non-sentential, and genuinely are assertions of truth-conditions, then we have strong evidence that the relation between an utterance and its intuitive truth-conditions can be mediated by processes quite unlike assignments to parts of a sentence and combination. If this phenomenon is genuine, then something like an existence proof has been provided for “free” contextual enrichment (cf. Carston (2002, p. 63)) of the sort discussed above. In this second section, I argue that the phenomenon is not genuine, or more precisely, that to describe it as “non-sentential assertion” is incorrect. Finally, in the last section, I consider a wide variety of apparent counterexamples to my proposal about the relation between utterances and their intuitive truth-conditions, and argue that in fact they involve sentences whose structure is more complex than it appears to be on the surface.

The purpose of “Context and Logical Form” was to introduce the general project, and defend some of the methodology used to establish hidden structure. Discussion of particular examples is therefore not pursued in great detail. The purpose of the second paper in this section, “On Quantifier Domain Restriction”, co-authored with Zoltan Gendler Szabo, was to investigate one case in great detail, that of quantifier domain restriction. When someone utters a quantified sentence, such as “Every bottle is in the fridge”, they usually intend to communicate something about a restricted domain of things. In this paper, Szabo and I argue that one can explain the relation between an utterance of a quantified sentence and its intuitive truth-conditions within the restrictions imposed by my proposal in “Context and Logical Form.” In particular, we argue that quantifier domain restriction is due to the presence of unpronounced structure in the structure of the quantified sentence uttered.

The hidden structure postulated in “On Quantifier Domain Restriction” was not entirely motivated by the arguments provided in that paper. For example, in that paper, following work by Kai von Stechow (1994), we argued that a domain variable is really a package of a function variable and an object variable.<sup>2</sup> The object variable is assigned objects, relative to a context, and the function variable is assigned a function from objects to properties (we also argue that domains are properly conceived of as properties rather

---

<sup>2</sup> Von Stechow associates these variables with the determiner rather than (as we do) with the noun.

than sets). So, in a sentence such as “Every fireman is tough”, there are really two domain variables, which together determine the quantifier domain for the quantified noun phrase “every fireman”. Relative to a context, the variable ranging over objects associated with “every fireman” is perhaps assigned a location, and the variable over functions is perhaps assigned a function from locations to the property of being an inhabitant in that location. Examples by von Stechow and Robin Cooper (such as “In every township, every fireman is tough”) show that the object variable can be bound by a higher quantifier, which provides evidence of a syntactic sort for the existence of that variable. But no one had provided similar evidence of a similar sort for the existence of the function variable that is involved in domain restriction. For instance, no one had provided examples in which the function variable associated with domain restriction can be bound by a higher quantifier, despite the fact that the function variable is required for the semantic analysis to work correctly. In “Nominal Restriction”, the third paper in this volume, I provide the missing motivation for the hidden structure postulated by Szabo and my analysis. I then argue that the existence of this structure explains not only the phenomenon of quantifier domain restriction, but also a number of other apparently distinct contextual effects.

One kind of apparently distinct contextual effect explained by the apparatus of quantifier domain restriction involves the context-sensitivity of adjectival constructions. As Hans Kamp pointed out many years ago, the sentence “John built a large snowman” is intuitively context-sensitive. Suppose that John has been injured recently, and as part of his therapy, has been building snowmen. On this particular occasion, he has built a snowman that is much larger than his previous snowmen. Then an utterance of “John built a large snowman” may express a truth, even though the snowman he built is not particularly large for a snowman built by a healthy person of his age. As it turns out, recognizing that nouns are associated with domain indices allows for a solution to this problem.

Much of “Nominal Restriction” concerns context-sensitivity associated with mass terms. As with count nouns such as “bottle” and “sailor”, mass nouns can occur in quantificational structures. So, when I declare that there is a little milk in the refrigerator, my utterance is not made true by the existence of a small drop of milk in the corner;

intuitively “little milk” ranges only over larger quantities of milk. The nominal restriction theory advocated by Szabo and me explains this phenomena; “milk” is associated with a domain index, restricting the quantification to quantities of a certain size. But it also explains other kinds of context-sensitivity associated with mass terms, context-sensitivity more typically explained by appeal to “loose use”. Many philosophers think that “water” denotes quantities of pure H<sub>2</sub>O. When we say that lakes and rivers contain water, we are speaking loosely (since the stuff that flows in lakes and rivers is filled with impurities). But note that count nouns do not allow a similar kind of loose use; we never describe housecats as impure tigers. This suggests to me that the idea that mass terms such as “water” and “gold” have, as their extensions, only quantities of pure chemical kinds, is mistaken. Rather, what we have here is a species of semantic context-sensitivity, rather than loose talk. In the final section of “Nominal Restriction”, I argue that the nominal restriction theory justifies this intuition; terms such as “water” and “gold” contain, in their extensions, quantities of impure substances. In contexts in which we speak of only H<sub>2</sub>O being water, we are employing domain restriction.

Many debates in the philosophy of language are confused by the fact that theorists employ distinct conceptions of the relation between semantics and pragmatics. In “Context and Logical Form”, I was very specific about my use of term “semantics”, and I distinguished it from the use of the term “semantics” employed by some other theorists, such as John Perry. But, under the urging of Jeffrey King, I soon came to realize that a more thorough discussion and justification of terminological choices was necessary. Out of our discussions of these matters came our co-authored paper, “Semantics, Pragmatics, and the Role of Semantic Content”, much of which is devoted to an overview of different conceptions of the semantics-pragmatics distinction.

The first part of the paper is devoted to justifying our favored way of drawing the semantics-pragmatics distinction. For example, some theorists choose to extrude from the semantic content of sentences relative to contexts any effects of speaker intentions. According to such theorists (most prominently Kent Bach, but also Emma Borg), the semantic content of “That is a table” is not a proposition, because the reference of “that” is determined by speaker intentions, and speaker intentions do not influence semantic content, even relative to a context. Presumably, one motivation for extruding the effects

of speaker intentions from semantic content is the aforementioned concern that adding context-sensitivity to the semantics would make the semantics unsystematic. We argue in this paper that there are certain clear ways in which not recognizing the semantic effects of extra-linguistic context (including perhaps speaker intentions) could make the semantics less systematic, by posing a problem for compositionality.

The final part of our joint paper is another defense of the central thesis of the essays in this book, namely that the scope of semantic content includes intuitive truth-conditions. King and I address a host of constructions that classically have been held to pose a threat to those who wish to tie intuitive truth-conditions closely to sentential content (conceived of as on my proposal). For example, consider the comparative “Eating some of the cake is better than eating all of it.” Intuitively, an utterance of this expresses the truth-condition that eating some but not all of the cake is better than eating all of the cake. But the result of composing the contents of the parts of the sentence “Eating some of the cake is better than eating all of it” does not appear to yield this intuitive truth-condition. Certainly, the content of “some of the cake” is not “some but not all of the cake” – one way to eat some of the cake is to eat all of it. Similarly, consider a conditional such as “If John marries Hannah and has kids, he will be happy, but if he has kids and marries Hannah, he won’t.” Intuitively, an utterance of this conditional has the intuitive truth-condition that if John marries Hannah and subsequently has kids, he will be happy, but if he has kids and subsequently marries Hannah, he won’t. But this content does not seem derivable from the contents of the words occurring in the conditional construction. In this paper, we argue that none of these classic problem cases poses worries for my proposal about the relation between utterances and their intuitive truth-conditions. The claim that these cases do raise worries for my proposal derives from an overly simplistic view of the semantics of terms such as “if...then” and “better than”.

The next paper in this volume is “Making it Articulated.” The central purpose of this paper is to show that the non-semantic explanations others have suggested to accommodate data about speakers’ intuitions about the truth-conditions of their utterances are uniformly unsuccessful. The worries with my proposed account of the relation between sentences in context and their intuitive truth-conditions all involve under-generation. That is, my proposal about the relation between utterances and their

intuitive truth-conditions is highly restrictive, and there are many examples of utterances that seem to have intuitive truth-conditions that cannot be captured by my proposal. In contrast, the processes that are appealed to in the pragmatics literature have the opposite problem; they tend to *over-generate*, predicting the availability of readings of certain utterances that simply are not there. My purpose in this paper is to emphasize the depth of the over-generation problem facing attempts to accommodate the data in other ways.

The way I bring out the over-generation problem facing pragmatic explanations of apparently semantic phenomena is by focusing on attempted pragmatic explanations of data that I account for by appeal to hidden syntactic structure. In particular, my focus is on “free enrichment” accounts of *bound* readings of certain sentences, such as “Every species has members that are old”, where the comparison class for “old” seems to be bound by the initial quantifier “every species”. In “Context and Logical Form”, I had argued that this example provides evidence for unpronounced structure associated with predicative adjective constructions. I had also (following others, such as Kai von Stechow) argued for a similar conclusion for the case of quantifier domains. Advocates of free pragmatic enrichment had responded to my arguments by appeal to free pragmatic enrichment with variables. My purpose in “Making it Articulated” is to bring out the severity of the over-generation problem facing advocates of pragmatic accounts of this kind of data, by criticizing free pragmatic enrichment with variables. As I argue in this paper, those who appeal to free pragmatic enrichment to explain away bound readings of sentences are at a loss to explain why many sentences lack readings they should have, if free pragmatic enrichment were an interpretive option.

Note that on my own proposal, what is said by an utterance of a sentence does require a good deal of what is traditionally called “pragmatics”. For on my view, each sentence contains a number of unpronounced context-sensitive constituents. For example, understanding what is said by an utterance of “Every bottle is in fridge” requires grasping the value of domain indices associated with “bottle”. This involves grasping the contents of both a function variable and an object variable. Just as one cannot “read off” the proposition expressed by an utterance of “That is green” from the sentence used, one cannot “read off” the proposition expressed by an utterance of “Every bottle is in the fridge” from the sentence used. The sentence used restricts the possibilities of

interpretation in very specific ways, thereby easing the interpretive problem facing the hearer. But the multiplicity of context-sensitive elements, both pronounced and unpronounced, nevertheless places demands on both speaker and hearer. The hearer must infer the values of the context-sensitive elements in the utterance she apprehends, or something close enough to them, and the cooperative speaker must ensure that her audience is in a position to grasp them, or something close enough to them.

So the difference between my proposal and proposals that postulate an interpretive process of free enrichment is *not* that the latter kinds of approaches require participants in a conversation to engage in a great deal of reasoning about the context to grasp what is intuitively expressed. My proposal also requires interlocutors to engage in reasoning about the context. Such reasoning is needed in order to resolve the values of context-sensitive elements in the sentence, and to make sure that one's hearers can easily resolve them. My proposal embodies a quite specific hypothesis about *how* sentences constrain the possibilities of their correct interpretations, one that is considerably more restrictive than proposals that postulate "free enrichment", yet promises to explain all the data without facing the threat of over-generating readings of sentences.

Chapter 6 of this volume is "Semantics in Context". I begin by setting out the challenge from context-sensitivity, and discussing some of the most difficult problem cases for a view such as my own, which seeks to reflect the intuitive truth-conditions of an utterance in the semantic content for the sentence uttered, relative to a context of use. One such difficult problem case is the problem of *deferred reference*. Consider, for example, an annoyed waitress muttering "The ham sandwich is irritating", or a doctor telling a nurse, "The kidney in room 103 needs a glucose drip." The intuitive truth-conditions of the first utterance are that the person who ordered the ham sandwich is irritating, and the intuitive truth-conditions of the second utterance are that the person with the kidney problem in room 103 needs a glucose drip. The problem with the phenomenon of deferred reference is that it does seem to affect the intuitive truth-conditions of an utterance, and yet incorporating it into the theory of meaning proper threatens the systematic nature of semantic theory (or so I maintain). In Sections I and II of the paper, I try to develop several principled reasons to think that the phenomenon of deferred reference is not semantic in nature. In so doing, I try to give criteria for when a

phenomenon that seems to affect intuitive truth-conditions should be incorporated into a semantic account of what is said, criteria that are consistent with the demand that a semantic theory must be responsive to facts about speaker intuitions about truth-conditions of utterances.

In the third, fourth, and fifth sections of the paper, I return to a discussion of one piece of methodology that I have used to establish the presence of hidden structure, namely the binding argument, which involves an inference from the existence of bound readings of a sentence to the existence of unpronounced variables in the logical form of the relevant construction. So from the fact that “Every species has members that are old”, I conclude that a predicative adjective construction such as “is old” involves an unpronounced variable associated with “old”. One way to reject this inference is to appeal to free pragmatic enrichment; the bound reading is explained by the fact that speakers have available to them an interpretive strategy that can add variables to sentences they hear. In “Making It Articulated”, I discuss and reject such accounts. My purpose in sections 4 and 5 of “Semantics in Context” is to evaluate two alternative accounts that also would block the inference from the existence of bound readings to the existence of unpronounced variables in logical form. Unlike free pragmatic enrichment, these alternative accounts are each semantic in nature. But I argue that neither suffices to account for the presence of bound readings in all cases.

In “Making it Articulated”, I level the charge of over-generation against advocates of free pragmatic enrichment. What I mean by “over-generation” in that paper is that such accounts over-generate *readings of sentences*; if free pragmatic enrichment were an interpretive option, there should be readings of sentences that simply are not interpretive options for utterances of those sentences. But there is a somewhat related over-generation concern facing my own proposal. It is not that my account predicts there to be certain non-existent readings of sentences; it is rather that, given the interpretive possibilities, my account requires too much hidden structure. The final section of “Semantics in Context” is devoted to countering this concern.

Since I began this project many years ago, a number of authors have published books developing their views. One of the most important contributions to this literature is

Francois Recanati's Literal Meaning. I include my review of this book, in order to allow the reader some sense of the relation between my views and opposing positions.

### Bibliography

- Almog, Joseph, Perry, John and Wettstein, Howard (ed.) (1989): Themes from Kaplan. Oxford: Oxford University Press.
- Bach, K.: 1994, 'Conversational Implicature', *Mind and Language* **9**, 124-162.
- Bezuidenhout, Anne (1997): "The Communication of De Re Thoughts", *Nous* **31.2**: 197-225.
- Borg, Emma 2004: Minimal Semantics. Oxford: Oxford University Press.
- Cappelen, Herman, and Lepore, Ernie (2004): Insensitive Semantics: in defense of semantic minimalism and speech act pluralism, Basil Blackwell, 2004.
- Cappelen, Herman, and Lepore, Ernie (2006): "Shared Content", in Oxford Handbook of Philosophy of Language, Oxford University Press, (eds.) E. Lepore and B. Smith (w. E. Lepore).
- Carston, Robyn (2002): Thoughts and Utterances: The Pragmatics of Explicit Communication (Oxford: Blackwell Press).
- Kai von Fintel (1994): Restrictions on Quantifier Domains, PhD Thesis, University of Massachusetts, Amherst.
- Gibbard, Allan (1981): "Two Recent Theories of Conditionals," in W. L. Harper, R. Stalnaker, and G. Pearce (eds), Ifs: 211-47.
- Grice, Paul (1989): Studies in the Way of Words (Cambridge: Harvard University Press).
- Grice, Paul (1989a): "Logic and Conversation", in Grice (1989): 22-40.
- Hawthorne, John (2004): Knowledge and Lotteries (Oxford: Oxford University Press).
- Heck, Richard (2002): "Do Demonstratives Have Senses?" *Philosophers' Imprint* **2**.
- Kaplan, David (1989): "Demonstratives", in Almog, Perry, and Wettstein: 481-565.
- Ludlow, Peter (1995): "Logical Form and the Hidden Indexical Theory: A Reply to Schiffer" *Journal of Philosophy* **92**: 102-107.
- Sperber, D. and Wilson, D.: 1986, Relevance, Harvard University Press, Cambridge.
- Stalnaker, Robert (1987): Inquiry (Cambridge: MIT Press).
- Strawson, Peter (1996): "On Referring", in The Philosophy of Language, Third Edition, A.P. Martinich, ed. (New York: Oxford University Press): 208-230.
- Timothy Williamson (2005): "Knowledge, Context, and the Agent's Point of View", in Preyer and Peter, ed., Contextualism in Philosophy: Knowledge, Meaning, and Truth (Oxford: Oxford University Press): 91-114.