A PROLEGOMENON TO SITUATION SEMANTICS

David J. Israel
Bolt Beranek and Newman Inc.
Cambridge, MA 02238

ABSTRACT

An attempt is made to prepare Computational Linguistics for Situation Semantics.

I INTRODUCTION

The editors of the AI Journal recently hit upon the nice notion of correspondents' columns. The basic idea was to solicit experts in various fields, both within and outside of Artificial Intelligence, to provide "guidance to important, interesting current literature" in their fields. For Philosophy, they made the happy choice of Dan Dennett; for natural language processing, the equally happy choice of Barbara Grosz. Each has so far contributed one column, and these early contributions overlap in one, and as it happens, only one, particular; to wit: Situation Semantics. Witness Dennett:

...situation semantics - [is] the hottest new topic in philosophical logic...[is] in some ways a successor or rival to Montague semantics.

And now Grosz:

In recent work, Barwise and Perry address the problem [of what information from the context of an utterance affects which aspects of interpretation and how?] in the context of a proposed model theory of natural language, one that appears to be more compatible with the needs of AI than previous theories.... [I]t is of interest to work in natural-language processing for the kind of compositional semantics it proposes, and the way in which it allows the contexts in which in an utterance is used to affect its interpretation.

What is all the fuss about? I want to address this question, but rather indirectly. I want to situate situation semantics in "conceptual space" and draw some comparisons and contrasts between it and accounts in the style of Richard Montague. To this end, a few preliminary points are in order.

A. The Present Situation

First, as to the state of the Situation Semantics literature. There is as yet no published piece of the scope and detail of either "English as a Formal Language" or "The Proper Treatment of Quantification in Ordinary English". Nor, of course, is there anything like that large body of work by philosophers and linguists - computational and otherwise - that has been produced from within the Montague paradigm. Montague's work was more or less the first of its kind. It excited, quite justifiably, an extraordinary amount of interest, and has already inspired a distinguished body of work, some of it from within AI and Computational Linguistics. The latter can hardly be said for

Situation Semantics (yet?).

So what is there? Besides a few published papers, each of them containing at least one position since abandoned, there is a book Situations and Attitudes literally on the very verge of publication. This contains the philosophical/theoretical background of the program - The Big Picture. It also contains a very brief treatment of a very simple fragment of ALIASS. And what, the reader may well ask is ALIASS? An Artificial Language for Illustrating Aspects of Situation Semantics, that's what. Moreover there is in the works a collaborative effort, to be called Situations and Discourse.^{*} This will contain a "Fragment of Situation Semantics", a treatment of an extended fragment of English. Last, for the moment, but not least, is a second book by Barwise and Perry, Situation Semantics, which will include a treatment of an even more extended fragment of English, together with a self-contained treatment of the technical, mathematical background. (By "self-contained". understand: not requiring either familiarity with or acceptance of The Big Picture presented in S&A.) The bottom line: there is very little of Situation Semantics presently available to the masses of hungry researchers.

B. Similarities

There are important points of similarity between Situation and Montague semantics, of course. One is that both are committed to formulating mathematically rigorous semantic accounts of English. To this end, both, of course, dip heavily into set theory. But this isn't saying a whole lot; for they deploy very different set theories. Montague, for a variety of technical reasons, was very fond of MKM, a very powerful theory, countenancing huge collections. MKM allows for both sets and (proper) classes, the latter being collections too big to be elements of other collections, and too big to be sets, say, of ZF. It also provides an unnervingly powerful comprehension axiom. B&P, on the other hand, have at least provisionally adopted KPU, a surprisingly weak set theory. Indeed, the vanilla version of KPU comes without an axiom of infinity and (more or less hence) has a model in the hereditarily finite sets. In that setting, even little infinite collections, like the universe of hereditarily finite sets, are proper classes, and beyond the pale. Enough for the moment of set theory, although we shall have to return to this strange land for one more brief visit.

More important, and perhaps more disheartening, similarities are immediately to hand. Both Montague and B&P - thus far - restrict themselves to the declarative fragment of English; Montague, for the obvious reason that he was a model theorist and a student of Tarski. For such types, the crucial notion to be explicated is that of "truth

^{*}The collaborators being B&P, Robin Cooper, Hans Kamp, and Stanley Peters.

of a sentence on an interpretation". Moreover, Montague showed no interest in the use(s) of language. Of course people working within his tradition are not debarred from doing so; but any such interest is an extra added attraction. The same point about model theory, broadly construed, holds for Barwise-Perry as well; they certainly aren't syntacticians. But in their case it is reinforced by philosophical considerations which point toward the use of language to convey information as the central use of language - hence, to asserting as the central kind of utterance or speech act. Thus, even when they narrow their sights to this one use, the notion that language is something to be put to various uses by humans to further certain of their purposes is not foreign to Situation Semantics.*

Second, B&P (again: so far) stop short at the awesome boundary of the period. Here again, this was only to be expected; and here again, the crucial question is whether their overall philosophical perspective so informs their account of natural language as to enable a more fruitful accommodation of work on various aspects of extended discourse. Barbara Grosz hints at a suspicion I share, that although at the moment much of what we have in this regard are promissory notes and wishful thinking, the answer is in the affirmative.**

II THE BIG PICTURE

The major point, however, concerns the primary focus of the work of Barwise and Perry as contrasted with that of Montague. Montague approached the problem of the semantics of natural language essentially as a model theorist, attempting to apply (newly) orthodox mathematical techniques to the solution of classical problems in the semantics of natural languages, many of which had to do with intensional contexts. After all, these new techniques - in the development of which Montague played a role - had precisely to do with the treatment of formal languages containing modal and other intensional constructions. What made a fragment of English of interest to Montague, then, was that it contained loads of such contexts. It is as if all of that wondrous machinery, and the technical brilliance to deploy it, were aimed at an analysis of the following sentence: While the temperature was rising, John seemed to be looking for a unicorn who was thinking about a centaur. What is astounding, of course, is that Montague should have been able to pull a systematic and rigorous treatment of such contexts out of the model-theoretic hat.

When we turn to Situation Semantics, on the other hand, we seem to be back in the linguistic world of first-grade readers: Spot ran. Dick saw spot run. Jane believed that Spot ran. Indeed, the major concern of Barwise-Perry is not the semantics of natural language at all. They have bigger (well, different) fish to fry. First and foremost, they are concerned with sketching an account of the place of meaning and mind in the universe, an account that finds the source of meaning in nomic regularities among kinds of events (situations), regularities which, in general, are independent of language and mind. For the frying of said fish, a treatment of cognitive attitudes is essential. Moreover, and not independently, for any attempt to apply their overall philosophical picture to the semantics of natural language, the propositional attitude contexts pose a crucial and seemingly

insuperable obstacle.*** Hence the fact that the book Situations and Attitudes precedes Situation Semantics - the first lays the philosophical foundations for the second. Thus the origin of their concern even with the classical problems of the propositional attitudes is different from, though by no means incompatible with, that of Montague's.

Something brief must now be said about this big picture. Here goes.

The work of B&P can be seen as part of a continuing debate in philosophy about the source of the intentionality of the mental - and the nature of meaning in general; a debate about the right account to give of the phenomenon of one thing or event or state-of-affairs being able to represent (carry information about) another thing or event or state-of-affairs. On one side stand those who see the phenomenon of intentionality as dependent on language - no representation without notation. This doctrine is the heart of current orthodoxy in both philosophy of mind and meta-theory of cognitive psychology. (See, by way of best example, [5].) It is also a doctrine widely thought to be presupposed by the whole endeavor of Artificial Intelligence.* On another side are those who see the representational power of language as itself based on the intentionality of mind.** The striking thing about Barwise and Perry is that, while they stand firmly with those who deny that meaning and intentionality essentially involve language, they reject the thesis that intentionality and meaning are essentially mental or mind-involving.

The source of meaning and intentionality is to be found, rather, in the existence of lawlike regularities - constraints - among kinds of events. For Barwise-Perry, the analysis of meaning begins with such facts as that: smoke means fire or those spots mean measles. The ground of such facts lies in the ways of the world; in the regularities between event types in virtue of which events of one type can carry information about events of other types. If semantics is the theory of meaning, then there is no pun intended in the application of semantic notions to situations in which there is no use of language and, indeed, in which there are no minds.

Meaning's natural home is the world, for meaning arises out of the regular relations that hold among situations, among bits of reality. We believe linguistic meaning should be seen within this general picture of a world teeming with meaning, a world full of information for organisms appropriately attuned to that meaning.***
[3]

There is yet another dimension to the philosophical debate, one to which Barwise and Perry often allude:

Some theories stress the power of language to classify minds, the mental significance of language, and treat the

***I shall return to this theme below.

*Who knows? Maybe it is.

**These latter can, in turn, be divided into those who seek a naturalistic, in principle physicalist, account and those who, like Frege and Church, pose no such demand.

***For an important philosophical predecessor, see [4].

*"A Fragment of Situation Semantics" will contain a treatment of certain kinds of English interrogatives; further out in the future, Situation Semantics will contain such a more extensive treatment.

**Breaking out of the straightjacket of the sentence is the job of Situations in Discourse.

classification of (external) events as derivative....A second approach is to focus on the external significance of language, on its connection with the described world rather than the describing mind. Sentences are classified not by the ideas they express, but by how they describe things to be...Frege adopted a third strategy. He postulated a third realm, a realm neither of ideas nor of worldly events, but of senses. Senses are the "philosopher's stone", the medium that coordinates all three elements in our equation: minds, words and objects. Minds grasp senses, words express them, and objects are referred to by them....One way of regarding the crucial notion of intension in possible world semantics is a development of Frege's notion of sense. [3]

Barwise and Perry clearly opt for the second approach. This is one reason for their concern with the problems posed by the propositional attitudes; for it has often been argued that these contexts doom any attempt at a theory of the second type. This is the burden of the dreaded "Slingshot" - a weapon we shall gaze at later. For the moment, though, I want simply to note the connection of this dimension with that about the source and nature of intentionality. Just as (some particular features of) a particular X-ray carries information about the individual on which the machine was trained, e.g., that its leg is broken, so too does an utterance by the doctor of the sentence "It's bone is broken", in a context in which that same individual is what's referred to by "it". One can, of course, learn things about the X-ray and the X-ray machine as well as about the poor patient; just so, one can learn things about the doctor from her utterance. In both cases, the gaining of this derivative information is grounded in certain regularities, in the one case mechanical, optical and electro-magnetic; in the other, perceptual, cognitive, and social-conventional. More to the point, in all cases the central locus of meaning is a relation, a regularity, between types of situation and the primary focus of significance is an external event or event-type.*

Now, alas, for that return to set theory. I have studiously avoided telling the reader what situations, events and/or event-types are. Indeed, I haven't even said which, if any of these are technical terms of Situation Semantics. Later I shall say enough (I hope) to generate an intuitive feel for situations; still, I have been speaking freely of the centrality of relations between events or between event-types. Set-theoretically-speaking, such relations are going to be (or be represented by) collections of ordered-pairs. Collections, but not sets. These collections are proper classes relative to KPU; so, if this be the last word on the matter, those very regularities so central to the account are not themselves available within the account - that is, they are not (represented by) set-theoretic constructs generated from the primitives by way of the resources of KPU. For all such constructs are finite.**

Needless to say, that isn't the last word on the matter. Still, this is scarcely the place for an extended treatment of the issue; I raise it here simply to drive home a point about that first

*Needless to say, we can talk about both minds and mental events and languages and linguistic events; the key point is simply that a language user is not "really" always talking first and foremost about his/her own mental state. We are not doomed to pathological self-involvement by being doomed to speak and think.

**Assuming that we stick to an interpretation within the hereditarily finite sets, as we can.

similarity between Montague and Situation Semantics. Montague wanted a very strong background theory within which models can be constructed precisely because he didn't want to have to worry about any (size) constraints on such models. B&P put their money on a very weak set theory precisely because they want there to be such constraints; in particular because they want to erect a certain kind of barrier to the infinite. Obviously, large issues loom on the horizon; let's leave them there.

I want now briefly to discuss 3 major aspects of Situation Semantics, aspects in which it differs fairly dramatically from Montague semantics. In passing, I will at least hint at the interrelationships among these. Aside from particular points of difference, remember that in the background there lurks a general conception of the use of language and its place in the overall scheme of things, a conception that is meant to inform and constrain detailed proposals.

III THE PRINCIPLE OF EFFICIENCY

One other respect in which Barwise and Perry are orthodox is their acceptance of a form of the Principle of Compositionality, the principle that the meaning of a complex expression is a function of the meanings of its constituents. This is the principle that is supposed to explain the productivity or generativity of languages, and the ability of finite creatures to master them. But for Barwise and Perry, an at least equally important principle is the Principle of the Efficiency of Language.* This principle is concerned with the ability of different people at different times and places and in different contexts to (re)use the self-same sentence to say different things - to impart different pieces of information. So, to adopt their favorite example, if Mitch now says to me, "You're dead wrong", what he says - what he asserts to be the case - is very different from what I would say if I were to utter the very same sentence directed at him.** The very same sentence is used, "with the same meaning"; but the message or information carried by its use differs. Moreover, the difference is systematically related to differences in the contexts in which the utterances are made.

Barwise and Perry take this phenomenon, often called indexicality or token-reflexivity and all too often localized to the occurrence of particular words (e.g., "I", "you", "here", "now", "this", "that"), to be of the essence of natural languages. They also note, however, that their relational account of meaning shows it to be a central feature of meaning in general.

[T]hat smoke pouring out of the the window over there means that that particular building is on fire. Now the specific situation, smoke pouring out of that very building at that very time, will never be repeated. The next time we see smoke pouring out a building, it will be a new situation, and so will in one sense mean something else. It will mean that the building in the new situation is on fire at the new time. Each of these specific smoky situations means something, that the building then and there is on fire. This is...event meaning. The meaningful situations had something in common, they were of a common type, smoke pouring out of a building, a type that means fire. This is ...event-type meaning...What a particular case of smoke pouring out of a building means, what it tells us about the

*B&P choose to call such principles "semantic universals" - an unhappy choice, I think.

**Which, of course, I would never do.

wider world, is determined by the meaning of smoke pouring out of a building and the particulars of this case of it. [3]

Moreover, B&P contend that the fact that modern formal semantics grew out of a concern with the language(s) of mathematics has caused those working within the orthodox model-theoretic tradition either to ignore or to slight this crucial feature.*

A preoccupation with the language of mathematics, and with the seemingly eternal nature of its sentences, led the founders of our field to neglect the efficiency of language. In our opinion this was a critical blunder, for efficiency lies at the very heart of meaning. [3]

A. A Little Background

Sure enough, indexicality gave nightmares to both Frege and Russell.** It might seem that the issue of indexicality did not escape Montague's attention; and it didn't. Indeed, as Thomason says, "As a formal discipline, the study of indexicals, owes much of its development to Montague and his students" [22]. (See especially [10] and [11, 12].) This last is most especially true with respect to the work of David Kaplan, both a student and a colleague of Montague's. For Kaplan disagreed with Montague precisely about the extent to which the formal treatment of contexts of utterances should be accommodated to the treatment of intensionality via possible worlds. And B&P start from where Kaplan leaves off. [7, 8]

I shall assume once again the right to be sketchy: Montague adopted a very narrow stance towards issues in pragmatics, concerning himself solely with indexicals and tense and not concerning himself at all with other issues about the purposes of speakers and hearers and the corresponding uses of sentences.*** In addition, the treatment of formal pragmatics was to follow the lead of formal semantics: the central notion to be investigated was that of truth of a sentence, but now relative to both an interpretation and a context of use or point of reference. (See [10, 11, 12, 18].) The "working hypothesis" was that one could and should give a thoroughly uniform treatment of indexicality within the model-theoretic framework deployed for the treatment of the indexical-free constructions. Thus, for example, in standard quantificational theory, one of the "parameters" of an interpretation is a domain or universe of discourse; in standard accounts of modal languages, another parameter is a set of possible worlds; in tense logics, a set of points of time. Why stop there? It is clear when we get to indexicals that the three parameters I've just mentioned aren't sufficient to determine a function to truth-values. Just think of two simultaneous utterances of "You are dead wrong" in the same world, with all other

things equal except speaker and addressee. In the interests of uniformity, stuff all such parameters into structures called points of reference, and who knows how many we'll need - see [9], where points of reference are called indices. Then the meaning of a sentence is a function from points of references into truth values.

A number of researchers working within the Montague tradition (in a sense there was no other) were unhappy with this particular result of Montague's quest for generality; the most important apostate being Kaplan.* There are complex technical issues involved in the apostasy, centrally those involving the interaction of indexical and intensional constructions - interactions which, at the very least, cast doubt on the doctrine that the intensions of expressions are total functions from the set of points of reference to extensions of the expression at that point of reference.** The end result, anyway, is the proposal for some type of a non-uniform two-step account. Montague's points of reference should be broken in two, with possible worlds (and possibly, moments of time) playing one role and contexts of use (possibly including moments of time) another, different, role.

In this scheme, sentences get associated with functions from contexts of use to propositions and these in turn are functions from contexts to truth-values. Contexts, upon "application" to utterances of sentences, yield determinate propositions; worlds (world-times) function rather as points of evaluation, yielding truth values of determinate propositions.***

B&P, however, go beyond Kaplan's treatment, and in more than one direction. Crucially, the treatment of indexicality proper is only one aspect of the account of efficiency, in some ways, the least intriguing of the lot. Still, to drive home the first point: as it is with smoke pouring out of buildings, so too is it with sentences. The syntactic and semantic rules of a language, conventional regularities or constraints, determine the meaning - the event-type meaning - of a sentence; features of the context of use of an utterance of that type get added in to determine what is actually said with that use. This is the event meaning of the utterance, also called its interpretation. Finally, that interpretation can be evaluated, either in a context which is essentially the same as the context of use, or some other; thereby yielding an evaluation of the utterance, (finally) a truth value.

B. Beyond Indexicality

For B&P, the features of the context of use go beyond those associated with the presence of explicit indexical items in the utterance - people with personal pronouns, places with "locatives", times with tense markers and temporal indicators. In particular they mention two such parameters: speaker connections and resource situations. Some aspects of the former can be looked on as aspects of indexicality, following the lines of Kaplan's treatment of demonstratives. But in other respects, e.g., the treatment of proper names, and certainly in the treatment of resource situations, the view they sketch seems to transcend the boundaries of even deviant model-theoretic semantics. For they mean to do justice, within a unified and systematic framework, both to the fact that the meaning of an utterance type

*Barbara Grosz hints at agreement with this judgment. "[O]ne place that situation semantics is more compatible with efforts in natural-language processing than previous approaches [is that] context and facts about the world participate at two points: (1) in interpretation, for determining such things as who the speaker is, the time of utterance.; (2) in evaluation, for determining such things as whether the relationships expressed in the utterance hold."

**For the former, see [14], see also [15].

***Stalnaker is a wonderful example of someone working within the Montague tradition who does take the wider issues of pragmatics to heart. See [19].

*See [7, 8]. Others included Stalnaker and Kamp. See [19, 20] and [6].

**The extension appropriate to sentences and clauses being truth values.

***There is even a version of this called "two-dimensional modal logic" [20].

"underdetermines" the interpretation of an utterance of that type and to the fact the interpretation of an utterance "underdetermines" the information that can be imparted by that utterance. It is a constraint they impose on themselves that they be able to account for significant regularities with respect to "the flow of information", in so far as that flow is mediated by the use of language and in cases where the information is not determined by a compositional semantic theory. And such cases are the norm. Compositionality holds only at the level of event-type or linguistic meaning. The claim is that seeing linguistic meaning as a special case of the relational nature of meaning - that meaning resides in regularities between kinds of situations - allows them to produce an account which satisfies this constraint.

C. Names

So, let me say something about proper names and something else about resource situations. Let us put aside for the moment the semantic type that poor little "David Israel" gets assigned in [13]. Instead, we shall pretend that it gets associated with some individual.* But which individual? Surely with one named "David Israel"; but there are bunches of such, and many, many more Davids. The problem, of course, is that proper names aren't proper.** Just as surely, at the level of linguistic meaning it makes no sense for me to get special treatment with respect to my name.*** Still, if you (or I) hear Mitch Marcus, right after my talk, complaining to someone that "David is dead wrong", we'll know who's being maligned. Why so? Because we are aware of the speaker's connections; more finely, of the relevant connections in this instance. At the level of event-type or linguistic meaning, the contribution of a name is to refer to an individual of that name.**** On the other hand, it is a feature of the context of use, that the speaker of an utterance containing that name is connected in certain ways to such and such individuals of that name. Surely Mitch knows lots of Davids and we might find him saying "David thinks that David is really dead wrong". Of course, he might be talking about someone inclined to harsh and "objective" self-criticism; probably not.

Just one more thing about names and speaker connections. I noted above that for B&P, the interpretation of an utterance event underdetermines the information carried by that event. The use of names is a locus of nice examples of this. It is no part of the interpretation (event meaning) of Mitch's complaint about me that my name is "David"; but someone who saw him say this while he (Mitch, that is) was surreptitiously looking my way can learn that my name is "David", or even that I am the David Israel who gave the talk on Situation Semantics. Even without that, someone could learn that Mitch knows (is connected with) at least one person so named. (Of course, there are possibilities for "misinformation" here, too.) Just so, when I

*Some possible individual? My grandmother, for one, would have disagreed. So, too, do B&P.

**Mostly not; but how about "Tristan Tzara", to pick a name out of a hat?

***English should have no truck with (even) benign analogues of bills of attainder.

****It's a nice question whether some names carry with them, at the linguistic level, "species" information as well. But surely it doesn't seem to be an abuse of English to call, say, a platypus "David Israel".

introduce myself by saying "I'm David Israel", the interpretation of what I say on that occasion is singularly uninteresting, being (roughly speaking) an instance of the law of self-identity. But I will have conveyed the information I wanted to, namely that I am a David Israel, that "David Israel" is my name (though not mine alone). That's why we engage in the (otherwise inexplicable) custom of making introductions. Anthropology aside, the central point is that Situation Semantics is meant to give us an account in which we can explain and predict such regularities in the flow of information as that exploited by the convention of introductions. This account must show how such regularities are related to the conventional regularities that determine the linguistic meaning of sentence types and the patterns of contextual determination which then generate the meanings of particular utterance events.

D. Definite Descriptions

An analogue of the problem of the impropriety of talk of proper names arises with respect to definite descriptions. Take a wild and woolly sentence such as "The dog is barking". Again, we want the denotations of such definite descriptions to be just plain individuals; but again, which individuals? Surely, there is more than one dog in the world; does the definite description fail to refer because of non-uniqueness? Hardly; at the level of sentence meaning, there is no question of it's referring to some one individual dog. Rather we must introduce into our semantic account a parameter for a set of resource situations. Suppose, for instance, that we have fixed a speaker, an audience and a (spatio-temporal) location of utterance of our sentence. These three are the main constituents of the parameter B&P call a discourse situation; note that this one parameter pretty much covers the contextual features Montague-Kaplan had in mind. Suppose also that a dog, otherwise unknown to our speaker and his/her audience, just walked by the front porch, on which our protagonists are sitting. When the speaker utters the sentence, he/she is exploiting a situation in which both speaker and audience saw a lone dog stroll by; he/she is not describing either that particular recent situation or such a situation-type - there may have been many such; the two of them often sit out on that porch, the neighborhood is full of dogs. Rather, the speaker is referring to a situation in which that dog is barking. Which dog? The one "contributed" by the resource situation; the one who just strolled by. It is an aspect of the linguistic meaning of a definite description that a resource situation should enter into the determination of its reference on a particular occasion of use; thus, an aspect of the meanings of sentences that a resource situation be a parameter in the determination of the interpretations (event meanings) of sentential utterances. Moreover, one can imagine cases where what is of interest is precisely some feature of which resource situation a speaker is exploiting on a particular occasion. And here, too, as in the case of names or, more generally, of speaker connections, the claim is that the relational theory of meaning and the consequent emphasis on the centrality of the Principle of Efficiency give Situation Semantics a handle on a range of regularities connecting uses of languages with varieties of information that can be conveyed by such uses.

IV LOGICAL FORM AND ENTAILMENT

As we have noted, Barwise and Perry's treatment of efficiency goes beyond indexicality and, as embedded within their overall account, goes well beyond a Kaplan-Montague theory. An important theme in this regard is the radical de-emphasizing of the role of entailment in their semantic theory and the correlative fixing on statements, not sentences, as the primary locus of interpretation. This is yet another way in which B&P go beyond Kaplan's forays beyond Montague.

I have said that in standard (or even mildly

deviant) model-theoretic accounts the key notion is that of truth on an interpretation, or in a model. Having said this, I might as well say that the key notion is that of entailment or logical consequence. A set of sentences S entails a sentence A iff there is no interpretation on which all of the sentences in S are true and A is false. From the purely model-theoretic point of view, this relation can be thought of as holding not between sentences, but between propositions (conceived of as the intensions or meanings of sentences). For instance, it might be taken to hold between sets of possible worlds. Still, it is presumed (to put it mildly) that an important set of such relations among non-linguistic objects have syntactic realizations in relations holding among sentences which express those propositions. Moreover, that sentences stand in these relations is a function of certain specifiable aspects of their syntactic type - their "logical form".

In artificial, logical languages, this presumption of syntactic realization can be made more or less good; and anyway, the connections between, on the one hand, syntactic types and modes of composition, and semantic values on the other, must be made completely explicit. In particular, one specifies a set of expressions as the logical constants of the language, specifies how to build up complex expressions by the use of those constants, operating ultimately on the "non-logical constants", and then - ipso facto - one has a perfectly usable and precise notion of logical form.

In the standard run of such artificial languages, sentences (that is: sentence types, there being no need for a notion of tokens) can be, and typically are, assigned truth-values as their semantic values. Such languages do not allow for indexicality; hence the talk about "eternal sentences". The linguistic meaning of such a sentence need not be distinguished from the proposition expressed by a particular use of it.* Once indexicality is taken seriously, one can no longer attribute truth-values to sentences. (Note how this way of putting things suggests just the unification of the treatment of indexicality with that of modality that appealed to Montague.) One can still, however, take as central the notion of a sentence being true in a context on an interpretation. The main reason for this move is that it allows one to develop a fairly standard notion of logical consequence or entailment at the level of sentences. Roughly, a set of sentences S entails a sentence A iff for every interpretation and for every context of use of that interpretation: if every sentence in S is true in a given context, then so too is A.

Barwise&Perry are prepared to deemphasize radically the notion of entailment among sentences. As they fully realize they must provide a new notion - a notion of one statement following from another.

At the very least then, our theory will seek to account for why the truth of certain statements follows from the truth of other statements. This move has several important consequences... There is a lot of information available from utterances that is simply missed in traditional accounts, accounts that ignore the relational aspect of meaning... A semantic theory must go far beyond traditional "patterns of inference"... A rather startling consequence of this is that there can be no syntactic counterpart, of the kind traditionally sought in proof theory and theories of logical form, to the semantic theory of consequence. For consequence is simply not a relation between purely syntactic elements.

*Hence part, at least, of the oddity of talk about using such a language by uttering sentences thereof.

What's at stake here? A whole lot, I fear. First, utterances - e.g., the makings of assertions - are actions. They are not linguistic items at all; they have no logical forms. Of course, they typically involve the production of linguistic tokens, which - by virtue of being of such and such types - may have such forms. (Typically, but not always - witness the shaking or nodding of a head, the winking of an eye, the pointing of a finger, all in appropriate contexts of use, of course.) Thus, entailment relations among statements (utterances) can't be cashed in directly in terms of relations holding among sentences in virtue of special aspects of their syntactic shape. Remember what was said above about the main reason for opting out of an account based on statements and for an account based on sentence(type)-in-a-context. If you don't remember, let me (and David Kaplan) remind you:

First, it is important to distinguish an utterance from a sentence-in-a-context. The former notion is from the theory of speech acts, the latter from semantics.* Utterances take time, and utterances of distinct sentences can not be simultaneous (i.e., in the same context). But in order to develop a logic of demonstratives it seems most natural to be able to evaluate several premisses and a conclusion all in the same context. [8]. (The emphasis by way of underlining is mine - D.I.)

A logic has to do with entailment and validity; these are the central semantic notions; sentences are their linguistic loci. This all sounds reasonable enough, except of course for that quite unmotivated presumption that contexts of use can't be spatio-temporally extended. And it seems correspondingly unreasonable when B&P opt out.

[T]he sentence "Socrates is speaking" does not follow from the sentences "Every philosopher is speaking", "Socrates is a philosopher" even though this argument has the same "logical form" (on most accounts of logical form) as ["4 is an integral multiple of 2", "All integral multiples of 2 are even" (so) "4 is even"]. In the first place, there is the matter of tense. At the very least the three sentences would have to be said at more or less the same time for the argument to be valid. Sentences are not true or false; only statements made with indicative sentences, utterances of certain kinds, are true or false. [3] (The example is mine - D.I.)

B&P simplify somewhat. It is not required that all three sentences be uttered simultaneously (by one speaker). Roughly speaking, what is required is that the (spatio)temporal locations of their utterance be close together and that the "sum" of their locations overlap with that of some utterance of Socrates. But that isn't all. The speaker must be connected throughout to one and the same individual Socrates, else a pragmatic analogue of the fallacy of equivocation will result. The same (or something similar) could be said about the noun phrase "every philosopher", for such phrases - just like definite descriptions - require for their interpretation a resource situation. One can imagine a case wherein a given speaker, over a specified time and at a specified place, connected to one and the same guy named Socrates, exploits two different resource situations contributing two different groups of philosophers, one for each of

*This is what is known in the trade as a stipulative definition.

the first two utterances. (The case is stronger, of course, if we substitute for the second sentence "Socrates is one of the philosophers.")

It must certainly seem that too much of the baby is being tossed out with the water; but there are alleged to be (compensating?) gains:

There is a lot of information available from utterances that is simply missed in traditional accounts, accounts that ignore the relational aspect of meaning. If someone comes up to me and says "Melanie saw a bear." I may learn not just that Melanie saw a bear, but also that the speaker is somehow connected to Melanie in a way that allows him to refer to her using "Melanie". And I learn that the speaker is somehow in a position to have information about what Melanie saw. A semantic theory must go far beyond traditional "patterns of inference" to account for the external significance of language... A semantic theory must account for how language fits into the general flow of information. The capturing of entailments between statements is just one aspect of a real theory of the information in an utterance. We think the relation theory of meaning provides the proper framework for such a theory. By looking at linguistic meaning as a relation between utterances and described situations, we can focus on the many coordinates that allow information to be extracted from utterances, information not only about the situation described, but also about the speaker and her place in the world. [3]

A. A Mild Anti-Climax?

Despite the heroic sentiments just expressed, B&P scarcely eschew sentences, a semantic account of which they are, after all, aiming to provide. In the formal account statements get represented by n-tuples (of course), one element of which is the sentence uttered; and if you like, it is the sentence-under-syntactic-analysis. (This last bit is misleading, but not terribly.) Other elements of the tuple are a discourse situation and set of speaker connections and resource situations. Anyway, there is the sentence. Given that, how about their logical forms?

Before touching on that issue, let me raise another and related feature of the account. This is the decision of B&P to let English sentences be the domain of their purely compositional semantic functions. For Montague, the "normal form" semantic interpretation of English went by way of a translation from English into some by now "fairly standard" logical language. (Such languages became fairly standard largely due to Montague's work.) Montague always claimed that this was merely a pedagogical and simplifying device; and he provides an abstract account of how a "direct" semantic interpretation would go. Still, his practice leaves one with the taste of a search for hidden logical forms of a familiar type underlying the grammatical forms of English sentences. No such intermediate logical language is forthcoming in Situation Semantics. First there is ALIASS:

An Artificial Language for Illustrating Aspects of Situation Semantics... has more of the structure of English than any other artificial language we know, but it does not pretend to be a fragment of English, or any sort of "logical form" of English. It is just what its name implies and nothing more.

Next, and centrally, there is English. The decision to present a semantic theory of English directly may make the end product look even more different than it is. It certainly has the effect

of depriving us of those familiar structures for which familiar "theorem provers" can be specified, and thus reinforces the sense of loss for seekers after a certain brand of entailments. Some may already feel the tell tale symptoms of withdrawal from an acute addiction.

There is, however, more to it than that - or maybe the attendant liberation is enough. For instance, are English quantifiers logical constants, and if so, which ones? Which English quantifiers correspond to which "formal" quantifiers? Is there really a sentential negation operator in English? Well, surely "It is not the case that" seems to qualify; but how about "not"? And how about conjunction?

Consider, for example, a statement made with the sentence (1) Joe admires Sarah and she admires him. Let us confine our attention to the utterances in which (1) has the antecedent relations indicated by (1') Joe-1 admires Sarah-2 and she-2 admires him-1. While sentence (1) is a conjunction of two sentences, a statement made with (1) in the way [with the connections - D.I.] indicated by (1') is not a conjunction of independent statements. [3]

In general, if u1 and u2 are two statements with the same discourse situations and connections (and resource situations?), some sense can be made out of a [sic] conjunctive or [sic] disjunctive statement, with u1 and u2 as "parts". But this is not true of arbitrary statements. Moreover, as in the case above, if we have a [sic] conjunctive statement, there may be no coherent decomposition of it into two independent statements. Talk of conjunctive and especially of disjunctive statements is likely to be wildly misleading. For the latter suggests, quite wrongly, that the utterer is either asserting one "disjunct" or the other. "A statement made using a disjunctive sentence is not the disjunction of two separate statements." ([3].)

In an appendix to "Situations and Attitudes", B&P suggest an analogue of propositional logic for statements within a very simple fragment of ALIASS. There is no (sentential) negation and no conditional; but more to the point, there are no unrestricted laws of statement entailment, e.g., between an arbitrary "conjunctive statement" and its two "conjuncts". Things get even worse when we add complex noun phrases to the fragment. The mind boggles.

V THE PROPOSITIONAL ATTITUDES

Here I shall be mercilessly brief.** The conventional wisdom, from Frege through to its logical culmination in Montague, has been that propositional attitude constructions are "referentially opaque"; more particularly, that substitution of co-designative singular terms within them does not preserve the truth-value of the whole. Within that orthodoxy there has been disagreement as to whether they are also hyperintensional; that is, as to whether

*See [1] passim; but especially the first two sections.

**Mostly because of the sheer "sex appeal" of the issues involved, and partly because of the availability of the relevant texts, it has been their treatment of the propositional attitude contexts that has made B&P a cause celebre among philosophers. This is unfortunate; so I intend to do my part, by somewhat underplaying this whole tangle.

substituting necessarily co-designative terms or logically equivalent sentences within them preserves truth-value. Montague himself thought they were not hyperintensional; but he countenanced the other view. (And sketched an account to handle it.) Barwise and Perry have the unique distinction of believing that said contexts are at least intensional and yet transparent to substitution of singular terms.* This position is both solitary and thought to be incoherent. If it were in fact untenable, that would be most unfortunate for them, as it is also more or less mandated by their adopting an approach centered on the external significance of language.

Indeed, there is supposed to be a proof that it is incoherent. The argument in question, which B.&P. call the slingshot, is sometimes supposed to show that all sentences with the same truth-value must designate the same thing; and hence, of course, that truth-values must be the primary semantic values of sentences. More usually and somewhat more technically, it has been supposed to show that if a sentential context allows substitution of logically equivalent sentences and co-designating definite descriptions salva veritate, then that context must be truth-functional. More clearly: that all modes of sentence composition are truth-functional unless they're opaque. That is, the only contribution made by a sentence, so embedded, to the whole can be its truth-value.

In fact, the slingshot is not a "knockdown proof"; that it is not is recognized by many of its major slingers(?). (See, for instance, [16, 17].) Instead, in all of its forms, it rests on some form or other of two critical assumptions:

1. logically equivalent sentences are intersubstitutable in all contexts salva veritate; or, such sentences have the same semantic value
2. the semantic value of a sentence is unchanged when a component singular term is replaced by another, co-referential singular term.

B&P reject the assumptions that underlie the slingshot. Here, too, especially with respect to the second assumption, tricky technical issues about the treatment of singular terms - both simple and complex - in a standard logic with identity are involved. B&P purposefully ignore these issues. They are interested in English, not in sentences of a standard logic with identity; and anyway, those very same issues actually get "transformed" into precisely the issues about singular terms they do discuss, issues having to do with the distinction between referential and attributive uses of (complex) singular terms. (See their discussion in [2] and chapter 7 of [3].) To show my strength of character, I'm not going to discuss the sexy issue of transparency to substitution of singular terms - except to say that, like Montague, B&P want a uniform treatment of singular terms as these occur both inside and outside of propositional attitude contexts; and that they also want to have it that the denotations of such terms are just plain individual objects. (How perverse!) Rather, I want to look briefly at the first assumption about

logical equivalence.**

A. The Relation Theory of Meaning

With respect to the end-result, what's crucial is that B&P reject the alleged central consequence of the slingshot: that the primary semantic value of a sentence is its truth-value. Of course, given what we have already said, a better way to put this is that for them, although statements are bearers of truth-values, the primary semantic value of a statement is not its truth value.

That honor is accorded to a collection of situations or events. Very roughly, the story goes like this: the syntactic and semantic rules of the language associate to each sentence type a type of situations or states-of-affairs; intuitively, the type actualizations of which would be accurately, though partially, described by any statement made using the sentence.* Thus:

Consider the sentence "I am sitting". Its meaning is, roughly, a relation that holds between an utterance *u* and a situation *a* just in case there is a (spatio-temporal) location *l* and an individual *a*, *a* is speaking at *l*, and in *a*, *a* is sitting at *l*.... The extension of this relation will be a large class of pairs of abstract situations. [3].

Now consider a particular utterance of that sentence, say by Mitch, at a specific location *l*'.

Then any situation that has [Mitch] sitting at *l*' will be an interpretation of the utterance. An utterance usually describes lots of different situations, or at any rate partially describes them. Because of this, it is sometimes useful to think of the interpretation as the class of such situations. Then we can say that the situations appearing in the interpretation of our utterance vary greatly in how much they constrain the world...When uttered on a specific occasion, our sentence constrains the described situation to be a certain way, to be like one of the situations in the interpretation. Or, one might say, it constrains the described situation to be one of the interpretations. [3]

B. On Logical Equivalence

If the primary semantic value of a sentence is a collection or a type of situations, then it is not surprising that logically equivalent sentences - sentences true in the same models - might not have the same semantic values, and hence, might not

**One point to make, though, is the following: the indexical personal pronouns are certainly singular terms. Frege's general line on the referential opacity of propositional attitude contexts certainly seems at its shakiest precisely in application to such pronouns - and in general to indexical elements. And remember if B&P are right, there is an element of "indexicality" in the use of proper names. If Mitch believes that David is dead wrong and I'm (that) David, then Mitch believes that I'm dead wrong. If Mitch believes that I'm dead wrong and I am David Israel, then Mitch believes that (this) David Israel is wrong. [14, 15]

*There is a class of exceptions to this, but I want not to get bogged down in details here.

*I should note that neither "situation" nor "event" is a technical term in Situation Semantics; though "event-type" is.

be intersubstitutable salva semantic value. Consider the two sentences: (1) Joe eats and (2) Joe eats, and Sarah sleeps or Sarah doesn't sleep. Let's grant that (1) and (2) are logically equivalent. But do they have the same "referent" or semantic value?

If we think that sentences stand for situations...then we will not be at all inclined to accept the first principle required in the slingshot. The two logically equivalent sentences just do not have the same subject matter, they do not describe situations involving the same objects and properties. The first sentence will stand for all the situations in which Joe eats, the second sentence for those situations in which Joe eats and Sarah sleeps plus those in which Joe eats and Sarah doesn't sleep. Sarah is present in all of these. Since she is not present in any of the situations that "Joe eats" stands for, these sentences, though logically equivalent, do not stand for the same entity. (Obviously B&P are here ignoring the "indexicality" inherent in proper uses of proper names - D.I.) [3]

Notice that without so much as a glance in the direction of a single propositional attitude context, we can see how B&P can avoid certain well-known troubles that plague the standard model-theoretic treatments of such constructions.* Moreover and most importantly, they gain these fine powers of discrimination among "meanings" without following either Frege into a third realm of sense or Fodor (?) deep into the recesses of the mind. The significance of sentences, even as they occur in propositional attitude contexts, is out into the surrounding world.**

VI THE BOTTOM LINE

What's the bottom line? Clearly, it's too soon to say. Indeed, I assume many of you will simply want to wait until you can look at least at some treatment of some fragment of English. Others would like as well to get some idea of how the project of Situation Semantics might be realized computationally. For instance, it is clear even from what little I've said that the semantic values of various kinds of expression types are going to be quite different from the norm and much thought will be needed to specify a formalism for representing and manipulating these representations adequately. Again, wouldn't it be nice to be told something at least about the metaphysics of Situation Semantics, about situations, abstract, actual, factual and real - all four types figure in some way in the account; about events, event-types, courses-of-events, schema, etc? Yes, it would be nice. Some, no doubt, were positively lusting after the scoop on how B&P handle the classic puzzles of intensionality with respect to singular terms. And so on. All in good time.

What I want to do, instead, is to end with a claim, Barbara Grosz's claim in fact, that

*On this point, compare, e.g., [22]. I do not mean to imply that there aren't good reasons for denying the hyperintensionality of the propositional attitudes. There are. See [21]. Still, no one doubts that such a position is counter-intuitive.

**Actually, there is another big issue looming here, the one that hangs on B&P's opting for a treatment which takes properties and relations, intensionally conceived, as primitive - instead, that is, of pretending that properties are functions from "possible worlds" into sets. Sets, of course, there are; but so too are there properties.

attention should be paid. At the moment, the bottom line with respect to Situation Semantics is not, I think, to be arrived at by toting up technical details, as bedazzling as these will doubtless be. Rather, it is to be gotten at by attention precisely to THE BIG PICTURE.

The relational theory of meaning, and more broadly, the centrality in Situation Semantics of the "flow of information" - the view that that part of this flow that is mediated by the uses of language should be seen as "part and parcel of the general flow of information that uses natural meaning" - allows reasoned hope for a theoretical framework within which work in pragmatics and the theory of speech acts, as well research in the theory of discourse, can find a proper place. In many of these areas, there is an abundance of insight, harvested from close descriptive analyses of a wide range of phenomena - a range hitherto hidden from both orthodox linguists and philosophers. There are now even glimmerings of regularities. But there has been no overarching theoretical structure within which to systematize these insights, and those scattered regularities, and through which to relate them to the results of syntactic and formal semantic analyses. Situation Semantics may help us in developing such a framework.

This last is a good point at which to stop; so I shall.

ACKNOWLEDGEMENTS

This research was supported in part by the Defense Advanced Research Projects Agency, monitored by ONR under Contract No. N00014-77-C-0378 and in part by the Office of Naval Research under Contract No. N00014-77-C-0371. Also, special thanks are due to B&P - who, of course, are solely responsible for all the weird ideas presented in this paper. Any remaining responsibility is to be charged to Mitch Marcus, who suggested I do this, and to Brian Smith, who agreed.

REFERENCES

- [1] Barwise, J. and Cooper, R.
Generalized Quantifiers and Natural Language.
Linguistics and Philosophy 4(2):159-219, 1981.
- [2] Barwise, K.J. and Perry, J.R.
Semantic Innocence and Uncompromising Situations.
In French, Vehling, and Wettstein (editors),
Midwest Studies in Philosophy, pages 387-404. University of Minnesota Press, Minneapolis, 1981.
- [3] Barwise, K.J. and Perry, J.R.
Situations and Attitudes.
Bradford Books, Cambridge MA, 1983.
- [4] Dretske, F.
Knowledge and the Flow of Information.
Bradford Books, Cambridge MA, 1981.
- [5] Fodor, J. A.
The Language of Thought.
Crowell, New York, 1975.
- [6] Kamp, H.
Formal Properties of 'Now'.
Theoria 37:227-273, 1971.

- [7] Kaplan, D.
Demonstratives.
1977.
unpublished manuscript.
- [8] Kaplan, D.
On the Logic of Demonstratives.
In French, Vehling, and Wettstein (editors),
Contemporary Perspectives in Philosophy
of Language, pages 401-412. University of
Minnesota Press, Minneapolis, 1979.
- [9] Lewis, D.
General Semantics.
In Davidson, D. and Harman, G. (editors),
Semantics of Natural Language, pages
169-218. Reidel, Boston, 1972.
2nd edition.
- [10] Montague, R.
Pragmatics.
In Thomason, R. (editor), Formal Philosophy,
pages 95-118. Yale University Press, New
Haven, 1974.
- [11] Montague, R.
Pragmatics and Intensional Logic.
In Thomason, R. (editor), Formal Philosophy,
pages 119-147. Yale University Press, New
Haven, 1974.
- [12] Montague, R.
Universal Grammar.
In Thomason, R. (editor), Formal Philosophy,
pages 222-246. Yale University Press, New
Haven, 1974.
- [13] Montague, R.
The Proper Treatment of Quantification in
Ordinary English.
In Thomason, R. (editor), Formal Philosophy,
pages 247-270. Yale University Press, New
Haven, 1974.
- [14] Perry, J.R.
Frege on Demonstratives.
Philosophical
Review LXXXVI(4):474-497, October, 1977.
- [15] Perry, J.R.
The Problem of the Essential Indexical.
Nous 13(1):3-21, 1979.
- [16] Quine, W.V.O.
Reference and Modality.
In From a Logical Point of View, pages
139-159. Harper & Row, New York, 1961.
2nd edition.
- [17] Quine, W.V.O.
Three Grades of Modal Involvement.
In The Ways of Paradox and Other Essays,
pages 156-174. Random House, New York,
1966.
- [18] Scott, D.
Advice on Modal Logic.
In Lambert, K. (editor), Philosophical
Problems in Logic, pages 143-173. Reidel,
Dordrecht, 1970.
- [19] Stalnaker, R.
Pragmatics.
In Davidson, D. and Harman, G. (editors),
Semantics of Natural Language, pages
380-397. Reidel, Boston, 1972.
2nd edition.
- [20] Stalnaker, R.
Assertion.
In Cole, P. (editor), Pragmatics, pages
315-332. Academic Press, New York, 1978.
- [21] Stalnaker, R.
Propositions.
1982.
unpublished ms.
- [22] Thomason, R.
Introduction.
In Thomason, R. (editor), Formal Philosophy,
pages 1-69. Yale University Press, New
Haven, 1974.