

Frege's Puzzle as some Problems in Science

Gabriel Segal

It is argued that the problem of explaining why one cannot intersubstitute co-extensive expressions within propositional attitude reports (PARs), *salva veritate*, should be addressed within the context of empirical enquiry. It is claimed that two main areas of enquiry will be involved: (i) an enquiry into why people make the judgements they do about entailment relations among PARs, (ii) an enquiry into the actual entailment relations among PARs. It is argued that (i) will itself divide into four distinct sub-projects, each concerned with a specific area of human cognition. Finally, it is suggested that (ii) should fall within the domain of a scientific psychological theory of propositional attitudes.

0. Introduction

Philosophical discussions of the semantics of propositional attitude reports have been dominated by Fregean puzzles of which the following is a typical example. "Archie Leach" is co-extensive with "Cary Grant" ("Archie Leach" being Cary Grant's original name) so if (1) is true, so is (2):

(1) Cary Grant is an actor

(2) Archie Leach is an actor.

But (3) doesn't entail (4):

(3) Fred believes that Cary Grant is an actor

(4) Fred believes that Archie Leach is an actor.

Why not? Certainly this is a puzzle. But who is it a puzzle for? What exactly is it that needs to be explained? What theoretical enterprise is responsible for explaining it? I propose that the puzzle (or versions of it) fall within two research projects: (a) explaining people's judgements about entailment relations among attitude reports, and (b) determining the actual entailment relations. I think that both of these – including (b) – are empirical projects, falling within different

domains of scientific psychology. In this paper I will sketch out a theory addressing (a), and offer a few comments about (b).

My claim is that when (a) and (b) are completed by their respective sciences, there will be no residual philosophical puzzle remaining. If that's right, then we have a pleasing confirmatory instance of a certain, not uncommon (but not universally accepted) picture of philosophy and its relation to science. The picture is that one important role of philosophy is to grapple with issues that are too abstract and nebulous to be addressed by empirical research projects and to clarify and focus those issues until they can be so addressed. I believe we are close at the turning point in respect of Frege's Puzzle. Thus what was once a problem for philosophers is now becoming a problem for scientists.

1. *Towards explaining people's judgements*

I'll argue that we need theories of four distinguishable areas of human cognition to explain people's judgements about propositional attitude reports – in particular, judgements of the sort illustrated above with respect to (1) and (2). The areas of cognition are: (i) Semantic competence; (ii) Implicit psychological competence; (iii) Explicit psychological competence; (iv) Pragmatic competence. I'll explain the role of each of these in turn, focusing mainly on (i).

1.1. *Semantic Competence*

I assume that peoples' understanding of sentences is owed, in part, to their unconscious knowledge of a semantic theory for their language. This is an empirical hypothesis along Chomskyan lines. I differ from Chomsky only in that I think that the semantic component of our unconscious knowledge of language is to be understood in terms of the classical semantic notions of truth, reference and satisfaction. (Chomsky himself seems not to favour this idea: see Chomsky (1995)). In fact I think that what we know is a Tarskian T-theory.¹ But that won't be particularly important in what follows. What is important is that the semantic theory attributes semantic properties to individual words and shows how to deduce sentence meaning from word meaning and syntactic structure.

1.1.1. Semantic theory and judgements of entailment

To what extent, in general, would a semantic theory account for judgements about entailment relations among sentences? Such jud-

gements will not always be explicable by semantics alone. If people judge that a sentence, S1 entails a sentence, S2, this might be accounted for by a combination of semantics with something else. Suppose someone judges that if S1, is true, then so must be S2. This might result from their holding (5)-(7):

- (5) S1 is true iff p
- (6) S2 is true iff q
- (7) If p, then q.

(5) and (6) would probably be accounted for by their semantic knowledge. (This would certainly be so if the semantic knowledge took the form of a T-theory – anyway, if not, (5) and (6) would be very easily derivable from the semantic theory). But (7) might well be accounted for by some other area of cognition. For example, if S1 and S2 are sentences about mathematics, then it would probably be the subject's mathematical knowledge that underlies their judgement of its truth.²

By contrast, undesirable entailments should be blocked by semantics, *ceteris paribus*. If people don't judge that a sentence S1 entails a sentence S2, then we shouldn't attribute to them knowledge of a semantic theory in which it is provable that if S1 is true, then S2 is true. For if we did, then we might have a hard time explaining why they refrained from making the judgement. Of course, there will be exceptional cases in which, due to complexity of the subject matter or neurosis or processing limitations, or some such, the subject doesn't draw out the consequences of her knowledge. The point is just that it's bad methodology for the semantic theorist to start by attributing knowledge that has certain consequences that the subject doesn't accept, and then look for excuses to spring the *ceteris paribus* clause.³

1.1.2. Some data

So, given that subjects do make the Frege judgements – that they do not e.g. think that (2) follows from (1) – we want to attribute to them knowledge of a semantic theory that also fails to ratify the entailments. This provides some valuable clues as to the nature of the semantic theory. For judgements about attitude reports can be sensitive to many properties of content sentences. In particular, they are sensitive to: choice of words, orthography, phonology and pragmatics. Please consider (8)-(11).⁴

- (8) a. Fred believes that all vixens are vixens
 b. Fred believes that all female foxes are vixens

Think of the case in which Fred believes that *vixens* applies only to adult female foxes. He is therefore inclined to say such things as "Not all female foxes are vixens". According to some people's judgements, we should take Fred at his words here, in spite of his slight misunderstanding, and so not attribute to him the belief that all female foxes are vixens. Thus they judge that (8a) could be true, while (8b) was false.

- (9) a. Fred believes that black and white are colours
 b. Fred believes that black and white are colors.

For (9), imagine the case in which Fred is a transatlantic traveller who happens to believe that the English word "colour" and the American word "color" differ slightly in meaning. He thinks that the former applies to all the colours, including black and white, but the latter only applies to, as it were, the colourful colours, and not to black or white. Here, parallel with the previous case, some would judge that (9a) would be true while (9b) was false.

- (10) a. Fred believes that Haahvahd is snotty
 b. Fred believes that Harverd is snotty.

(10) is analogous: Fred is a rural lad from Wisconsin, where he has heard of an important university, the name of which is pronounced "Harverd". He visits Cambridge, MA, and comes across a liberal arts college which the locals call "Haahvahd". He doesn't realise that Harverd is Haahvahd. In such a case, according to some, (10a) might be true while (10b) was false.

- (11) a. Fred said he was really impressed
 b. Fred said he was really impressed

In respect of (11): Fred attends a lecture, which he takes to be a complete fiasco. On leaving, he says, sarcastically, "I was really impressed". According to certain intuitions, I could report truly on Fred's comment by saying "Fred said he was really impressed", if I mimicked his sarcasm when I said "he was really impressed". Were I to omit the sarcasm, however, and say it straight, then my report would be false.

I emphasise that it is only many speakers who make these jud-

gements, not all. But a semantic theory should account for everyone, including those many. Specifically, since those who make the judgements might be rational beings and competent speakers, we should suppose that the meanings of the sentences in each of the (a)/(b) pairs differ: the meaning of (8a) must be distinct from the meaning of (8b), and so on. If not, if (8a) and (8b) were identical in meaning, then how could a rational, competent speaker envisage a situation in which (8a) was true and (8b) false? So we should attribute knowledge of a semantic theory that assigns appropriately different semantic properties to the relevant sentences.

1.1.3. Sketch of a semantic theory

In order to achieve this goal, I propose that we build all the relevant properties into the semantic objects of attitude reports. I suggest that the utterance of the content sentence refers to what I call its own "Total Form".⁵ A total form is a complex linguistic object with syntactic, semantic, phonological and pragmatic properties. Thus (12) partially and schematically tells us about the total form of some specific utterance of "Cary Grant acts" ("TF(ut("s"))" means roughly "the total form of the utterance of string s"):

- (12) TF(ut("Cary Grant acts")) = [S [NP [N Cary Grant]] [VP [V acts]]]

where: Semantic properties are blah blah, phonological properties are so and so, pragmatic properties are such and such. Think of total form as a syntactic tree, with relevant orthographic, semantic etc. properties associated with the appropriate nodes.

Given the right axioms in the T-theory, it is possible to generate T-theorems (very roughly) of the form (13):

- (13) An utterance of "a believes that S" is true iff a believes TF(ut("S"))

There is an odd feature of this account, one shared by any quotational theory and by Davidson's (1968) theory that the objects of the attitude reports are token utterances. On these views what believers are said to believe are linguistic objects in the reporter's language. If the reportee speaks no language, or does not share the reporter's language, this seems a bit strange. Thus "Thales believed that everything is (probably) made of water" is true (I think). But how can Thales have believed the TF of an English sentence that he has never heard and would not have understood?

One way to address the problem would be to adapt Davidson's (1968) 'samesaying' idea. A slightly neater one involves the notion of 'expression'. Think of it this way. A total form is a representation, like a sentence or an utterance. Since this is so, total forms can be used to express what people believe, desire and so on. Thales had a belief that he could have (and presumably did) express using a sentence of his language. I can express that same belief using a total form of my language. So we can understand it this way: 'Thales had a belief that is expressed by TF(ut("Everything is (probably) water"))'.

Note that this talk of 'expression' is not part of the semantic theory itself. It is just an intuitive explanation of the way 'believes' works in the semantic theory. The language faculty represents 'believes' as expressing a two-place relation that holds between individuals and TFs. And we can get an intuitive handle on the nature of the relation by using the notion of 'expression'.⁶

1.1.1.4. An apparent problem

On the TF theory, terms within content sentences play a role in determining which TF the content sentence refers to. This explains away the apparent failure of intersubstitutivity, in an obvious way. But, since TFs are so fine-grained, it blocks all entailment relations among PARs the content sentences of which differ in any phonological, syntactic, semantic or pragmatic properties. But this now creates an apparent problem. Often people do accept entailments of the form: "a believes that S" is true, so "a believes that S'" is true, where S differs from S'. Relatedly, many content sentences are appropriate for expressing any given belief. In fact we are often very liberal in such judgements.

Consider the following, (14):

- (14) a. Sid (on Sunday): "Hattie will be here in exactly a week"
 b. Eric (the next Saturday): "Sid thinks Hattie will be here tomorrow".

We often report beliefs this way. Sometimes this is because we expect people to believe obvious consequences of other things they believe. So it may be that Eric thinks that Sid himself will make the relevant inference. But sometimes – often, indeed – this is because we take many rather different sentences all to be appropriate for expressing a given belief. A similar phenomenon is clear also in indirect speech. Thus (15):

- (15) a. Sid (on Sunday): "Hattie will be here in exactly a week"
 b. Eric (the next Saturday): "Sid said Hattie will be here tomorrow".

In this case it's not that Eric thinks that Sid says obvious consequences of other things he's said. It's just that, in the context, he takes *Hattie will be here tomorrow* to be appropriate for expressing Sid's original speech act. Given this, it's presumably also true that, in this context, Eric would judge (16) to be valid:

- (16) If "Last Sunday, Sid said that Hattie will be here in exactly a week" is true, then "Sid said that Hattie will here tomorrow" is true.

The complete picture is thus that while in certain contexts many people are very strict about entailments among PARs (illustrated by (8)-(11)), it's also true that in many contexts many people are very liberal about them. We need a theory that can account for everyone in all contexts. This is where the other areas of cognition come in.

1.2. Implicit Psychological Competence

People who understand PARs grasp the concepts of propositional attitude psychology. That is to say, they have an understanding of beliefs, desires, perceptions, inferences and so on. Let us call this understanding the 'Psychology Faculty'. It is clear that, whatever the nature of the Psychology Faculty, it has a role to play in our judgements about entailment relations among PARs. To illustrate, suppose someone judges that if "Fred believes that all vixens are vixens" is true, so must be "Fred believes that all female foxes are vixens". (17) might tell a crucial part of the story (where 'KoS' stands for knowledge of semantics, 'PF' for the Psychology Faculty and 'F(Fred)' stands for a representation in the Psychology Faculty):

- (17) a. "Fred believes that all vixens are vixens" is true iff
 Fred believes TF(ut("All vixens are vixens")) [KoS]
 b. "Fred believes that all female foxes are vixens" is true
 iff Fred believes TF(ut("All female foxes are vixens"))
 [KoS]
 c. Fred believes TF(ut("All vixens are vixens")) iff
 F(Fred) [Interface of KoS with PF]
 d. Fred believes TF(ut("All female foxes are vixens")) iff
 F(Fred) [Interface of KoS with PF]
 e. "Fred believes that all vixens are vixens" is true iff
 "Fred believes that all female foxes are vixens" is true.
 [Inference from (i)-(iv)].

In effect, the idea is just that the subject has one way of representing Fred's belief that all vixens are vixens in thought and two ways of expressing this in language.

What is this 'Psychology Faculty'? There is good evidence that it is like the Language Faculty in being largely innate. It has a very specific pattern of development that appears to be universal across the species. The development is very rapid, the essentials normally being in place by 4.5 years of age. And its acquisition exhibits the standard poverty of stimulus problem that argues for innateness. Further, there is a genetic deficit that is more or less specific to understanding of psychology: autism. High functioning autistics, to a good approximation, can be good at almost any area of thought except psychology, which they are unable to understand. Moreover, there is a genetic deficit that goes the other way: Williams Syndrome. To a reasonable approximation, Williams Syndrome subjects are good at language and at psychology, but seriously retarded in all other respects.⁷

Although the Psychology Faculty has been studied a great deal by experimental and developmental psychologists, not very much is known about the actual nature of the representations it deploys: not very much is known about the details of 'F(Fred)'. One proposal, due to Alan Leslie (1987), is structurally very similar to the TF proposal. Leslie's idea is that the Psychology Faculty (which he calls 'ToMM' – theory of mind mechanism) represents attitudes by quoting sentences in the language of thought (LOT). So Fred's believing that all vixens are vixens would be represented by the LOT translation of "Fred believes 'all vixens are vixens'". Another, quite different proposal from Josef Perner (1991) is that attitudes are represented by mental models of the kind discussed by Johnson-Laird (1983).

Neither of these proposals has been developed in depth or detail, and both are, at present, sketches rather than developed theories. This is an area that needs to be studied much more, understood much better, if we are to solve Frege's Puzzle in all its manifestations. For the Puzzle arises here once again. On the view I've proposed, our thoughts about attitudes are really separate from our linguistic representation of attitudes – or, more exactly, of our Language Faculty's representations of sentences attributing attitudes.⁸ So, on this view, we have non-linguistic representations of the attitudes and the Puzzle recurs immediately. Both Leslie's and Perner's proposals are designed to address the Puzzle.

A point to note here is that the constraints on solving the Puzzle in this context – the context of studying the Psychology Faculty – are different from those on solving it in natural language semantics (here

conceived as the study as the semantic component of the Language Faculty). The account of representations in the PF is answerable to data about acquisition and deficits that is quite specific to the PF and independent of analogous data relevant to the Language Faculty. For example, it needs to be explained why autistics can understand linguistic and other forms of representation without understanding propositional attitudes. Moreover, as we will see in a moment, pragmatics has an important role to play in explaining data about talk about propositional attitudes, while it seems unlikely that it can play a similar role in the explanation of the Psychology Faculty.

1.3. *Explicit Psychological Competence*

If, as I have claimed, the PF is universal across the species, then something further needs to be posited to explain individual differences in respect of judgements about propositional attitudes in general and entailment relations among PARs in particular. One crucial component of this something further is the more or less explicit collection of views one acquires as one matures. These can become particularly prominent if the one spends time in a psychology or philosophy department in a university.

An explicit theory of beliefs may affect one's intuitions about what follows from what and what counts as an acceptable report. I'll illustrate this with a real case. I, myself, believe that if someone believes that a fortnight is a fortnight, then they must believe that a fortnight is fourteen days. But Tyler Burge (1978) doesn't accept this inference. This is because Burge holds that one can deploy the concept of a fortnight even if one doesn't fully grasp it: e.g. one can believe that a fortnight is ten days. Burge holds this, because he holds that one can possess a concept by virtue of having the right relations to the word. Thus: a subject who has the expression "a fortnight" in her repertoire, and has a reasonable (though imperfect) grasp of its meaning, and who is disposed to regard her use of the word as open to correction by experts, could, in virtue of these facts, possess the concept of a fortnight. And this concept could feature in the beliefs that a fortnight is ten days, and that a fortnight is not fourteen days.

I don't accept this theory of concept possession. Rather, I think that if someone says "A fortnight is ten days" then – assuming they understand "is ten days", are being sincere and so on – their word "fortnight" does not express the concept (or, better, my concept) of a fortnight. So, on my view it would be wrong to report their speech act or the belief behind it using the word "fortnight". Whether my view of concept possession is preferable to Burge's is not relevant here. What

is relevant is that our differing views lead to different judgements about PARs. For this indicates that a complete theory of people's judgements about PARs will need to take into account the role of explicit psychological theory as well as the PF.

1.4. Pragmatics

There is a great deal of pragmatic flexibility to attitude reporting (I suppose that this is widely agreed). To illustrate, let's return to Sid, Eric and Hattie from example (15), repeated below:

- (15) a. Sid (on Sunday): "Hattie will be here in exactly a week"
 b. Eric (the next Saturday): "Sid said Hattie will be here tomorrow"

In a standard sort of context, Eric's report is true. But in other, more restrictive contexts, it would count as false. Suppose (to pick a rather contrived example) that Sid is forever saying that Hattie will be here tomorrow. Think of Sid as like a character from a Samuel Becket play: each day he has the conviction that Hattie will appear the day after. Or think of him like a certain Gaulish chief, who thinks each that the sky will fall on his head tomorrow – although, as he also thinks, tomorrow never comes. However, on the Saturday in question, Sid undergoes a major change of outlook, and, on that particular day, he refrains from saying "Hattie will be here tomorrow". In this case Eric's report might well be judged false.

In this latter case, let's suppose, Eric is specifically talking about Sid's psyche, and Sid's *idée fixe* about Hattie's imminent arrival has been under explicit discussion. In such a context the precise formulation of Eric's report is important. The exact choice of words may not matter. But the report must preserve something like the Character of Sid's original utterance, in David Kaplan's (1989) sense.

We can also imagine that on this same Saturday, the one on which Sid undergoes his change of outlook, Eric is involved in a quite different conversation. This conversation is focused not at all on the eccentricities of Sid's psyche, but concerns in a general way who said what about when Hattie was due. In this context Eric's report (15b) is in order.

This case illustrates how, in one and the same worldly situation, the truth value of a speech report can vary with the interests of the speakers. This seems to be a pragmatic matter. The interests of the speaker and audience determine which features of the TF are relevant. As we saw above, in a very strict context, some might judge that the precise wording, spelling or pronunciation might be relevant.

In a slack context, we only care about the gist: any TF which is roughly the same in content (on some rather coarse notion of content) counts as expressing the reportee's attitude.

Pragmatics thus plays a crucial role. For it determines the extension of the relations expressed by propositional attitude verbs in specific contexts of utterance. If pragmatic factors determine which features of TFs are relevant in context to their appropriateness for expressing a reportee's attitude, then the pragmatic factors determine, in context, which reportee/TF pairs stand in the relations expressed by "says", "believes" and so on. This sort of pragmatic determination of the extension of a term doesn't seem to be unusual. Compare, e.g. "tall", "is a good musician" and "is spicy".

2. Discussion

Here is a rough sketch of how it might all hang together. The reporter – Eric, let's say – forms a view about an attitude of a reportee – say about a belief of Hattie's. This means he has a representation of Hattie's belief in his implicit or explicit psychological theory. (For simplicity's sake I will ignore questions about the relationship between implicit an explicit psychological competence).

Suppose that Eric thinks that Hattie believes that the content of mental states is locally supervenient on the physical micro structure of subjects. He wishes to report on this belief. Language production mechanisms – those mechanisms responsible for finding words to express oneself, mechanisms the functioning of which remain largely mysterious – need to select a TF which, in the conversational context, expresses Hattie's belief. Eric deploys his knowledge of pragmatics to select the appropriate level of grain (whether there's a role for marking distinctions by phonology, choice of words among synonyms etc.). He also takes into account his knowledge of the audience's knowledge of philosophy, their acquaintance with the jargon and so on. All this is pragmatics. With these matters in hand, somehow or other, a particular TF – TF* – is selected from among those made available by the Language Faculty. He then says: "Hattie believes that s", where "s" is the string of words of which TF* is the TF. His report is then true if his various faculties have done their jobs correctly.

3. Scientific Psychology

What ordinary humans think about attitudes is one thing. The truth may be another. It is the business of a scientific psychology of

the attitudes to find out what they are really like. I envisage a developed science of the attitudes deploying its own technical language. Like any science, psychology should be clear about how it individuates elements of its subject matter and what the logical properties of its technical representations are. So it should be clear when two representations are equivalent ways of describing a single type of state, and it should be clear when one representation follows logically from another. So Frege's Puzzle must, in a sense, be addressed also within real psychology.

But here again the constraints on the solution are particular to the case. For one thing, the psychologist need not be specially concerned about lay intuitions. But, more important, the psychologist can draw on features of the attitudes that aren't specified by common-sense psychology. For example, if it turns out that Jerry Fodor is right about LOT, then it will be in order to refer to LOT sentences when describing attitudes. So the psychologist might use a representation a somewhat like (17) to express the claim that Olly believes that Stan is confused:

- (17) Olly stands in computational relation BEL to an LOT sentence which means that Stan is confused.

Or maybe the attitude report would describe the LOT sentence in detail.

I'll refrain from further speculation about the language of future psychology. The relevant future may be a long way off, and the speculations are somewhat idle. The point, though, should be clear enough: the version of Frege's Puzzle that arises for psychology is just one aspect of that science's empirical problematic.

So, to conclude: what was once a problem in philosophical logic is now roughly five problems in different branches of empirical psychology.

Address of the author

Department of Philosophy, King's College, London, e-mail: duddy037@bay.cc.kcl.ac.uk

Notes

- ¹ The conception of semantics adopted here is defended in detail in Larson & Segal (1995). See chapter one for a defence of the Chomskyan approach to linguistics.

stic competence in general and chapter two for a defence of the idea – much influenced by Donald Davidson (1984) – that knowledge of a T-theory accounts for semantic competence. Much of what follows is drawn from chapter 11.

² The conditional in (7) could be of any strength, from material to something with strong modal force. The actual strength will probably vary from case to case, depending on the subject matter of (7) and the area of cognitive competence that deals with it.

³ Semantic theory and judgements of inferential relations are discussed in chapters 2, 11, and 12 of Larson & Segal (1995), where the points of the previous paragraphs are made.

⁴ Examples like (8) have been around a long time. They are often used to show that it's not just failures of co-extensiveness that give rise to Frege puzzles: if any expressions are synonymous, then "female fox" and "vixen" are. (9) is inspired an example of Nathan Salmon's, involving "Catsup" and "Ketchup", (10) is an example of Peter Ludlow's, (11) is from Larson & Segal (1995) ch. 11.

⁵ This idea is sketched at the end of Larson & Segal (1995) ch. 11. It is an immediate descendant of the 'Interpreted Logical Form' theories of Harman (1972), Higginbotham (1986), Segal (1989) and Larson & Ludlow (1993). I think of the idea as very close in spirit to Davidson's (1968) theory that the objects of attitude reports are token utterances. The term 'total form' was suggested to me by Richard Larson.

⁶ Analogously, remarks about 'samesaying' are best not construed as part of the semantic theory. For discussion see Segal (1989), Burge (1986), Larson and Ludlow (1993), Larson and Segal (1995).

⁷ I argue for the genetic basis of the Psychology Faculty in Segal (1995).

⁸ Separate, not necessarily ontogenetically independent. It may be that the development of the PF somehow depends on the acquisition of language. The claim is only that once developed, the PF and the LF deploy different representations.

References

- ALMOG, JOSEPH, JOHN PERRY & HOWARD WETTSTEIN, eds. (1989), *Themes from Kaplan*, Oxford: University Press.
- BURGE, TYLER (1978), "Belief and Synonymy", *The Journal of Philosophy* 75:119-138.
- BURGE, TYLER (1986), "On Davidson's 'Saying That'", in LePore 1986:190-209.
- CARRUTHERS, PETER & PETER SMITH (eds.) (1995), *Theories of Theories of Mind*, Cambridge, Cambridge University Press, 1995.
- CHOMSKY, NOAM (1995), "Language and Nature", *Mind* 104:1-61.
- DAVIDSON, DONALD (1968), "On Saying That", *Synthese* 19:130-46, reprinted in DAVIDSON 1984.
- DAVIDSON, DONALD (1984), *Inquiries into Truth & Interpretation*, Oxford, Clarendon Press.
- HARMAN, GILBERT (1972), "Logical Form", *Foundations of Language* 9:38:65
- HIGGINBOTHAM, JAMES (1986), "Linguistic Theory and Davidson's Program in Semantics", in LePore 1986:29-45.
- JOHNSON-LAIRD, PHILIP (1983), *Mental Models*, Cambridge, Cambridge University Press.
- KAPLAN, DAVID (1989), "Demonstratives", in Almog *et al.* (1989).

- LARSON, RICHARD & PETER LUDLOW (1993), "Interpreted Logical Forms", *Synthese* 95:305-55.
- LARSON, RICHARD & GABRIEL SEGAL (1995), *Knowledge of Meaning: An Introduction to Semantic Theory*, Cambridge, MA, the MIT Press.
- LEPORE, ERNEST, ed. (1986), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Oxford, Basil Blackwell.
- LESLIE, ALAN (1987), "Pretense and Representation: The origins of 'theory of mind'", *Psychological Review* 94:412-426.
- PERNER, JOSEPH (1991), *Understanding the Representational Mind*, Cambridge MA, MIT Press.
- SEGAL, GABRIEL (1989), "A Preference for Sense and Reference", *The Journal of Philosophy* 86:73-89.
- SEGAL, GABRIEL (1995), "The Modularity of Theory of Mind", in Carruthers & Smith (1995:141-57).