

How should we revise the paratactic theory? *

Keith Frankish

1.

Donald Davidson's original 1969 statement of the paratactic theory of indirect discourse (PT) involved two claims. First, there was a claim about logical form:

LF: The logical form of an indirect speech report of the form 'A said that p ' is the paratactic pair:
A said that. $>u$

where u is an utterance, 'that' a demonstrative referring to u (the arrowhead here indicating demonstration), and 'said' a two-place relation (henceforth written 'samesaid'¹) between persons and utterances, such that 'A samesaid u ' is true if and only if A made an utterance with the same import as u .

Davidson also added a claim about the surface form of indirect speech:

SF: The surface form of an indirect speech report coincides with its logical form.

That is to say, the surface form of indirect speech is itself paratactic and requires only the insertion of a period ('a tiny orthographic change ... without semantic significance', as Davidson put it) to be made explicit. If SF is true, then the 'that' of 'says that' is, grammatically speaking, a demonstrative, and the utterance to which it refers is just the reporter's utterance of the words immediately following it.²

Early objectors to PT tended to attack LF directly, arguing that it did not yield a satisfactory account of the truth conditions of indirect speech reports. Defenders of PT usually found ways of answering, or at least of dodging, these objections and the theory survived and prospered. More recent critics, however, have concentrated their fire upon SF, and this flanking attack may ultimately prove more damaging for PT.

The case against SF is strong. In an elegant 1986 paper, Segal and Speas marshalled a range of evidence, all tending to show that, from a grammatical point of view, the 'that' of 'says that' is a complementizer, not a demonstrative (Segal and Speas 1986). The inadequacy of SF is also revealed in its treatment of certain ambiguous reports, as Ian Rumfitt has recently pointed out. Suppose someone utters the sentence

* This is the author's version of an article published in *Analysis* 56(4): 251-63, 1996. It differs in a few minor ways from the published version.

¹ 'Samesaid' is silently substituted for 'said' in quotation, throughout.

² I shall call these words the *content clause*, and the words that introduce them the *reporting clause*.

(1) John said that someone heard the shooting of the hunters.

(The example is Rumfitt's.) Without a clarifying context, this utterance is ambiguous. (It isn't clear whether the speaker is claiming that John said someone heard the hunters shooting or that John said that someone heard the hunters being shot.) Moreover, since the ambiguity is structural, rather than lexical, it ought to be possible to construct two distinct logical forms for (1), corresponding to the two possible readings of it. But if the logical form of (1) is to be obtained simply by inserting a period, then we will be able to construct just one such form, namely

(2) John samesaid that. >*u*

where *u* denotes the reporter's utterance of the sentence 'Someone heard the shooting of the hunters'. Nor does it seem likely that further analysis of (2) will reveal the required duality of forms. For (2) is itself unambiguous. It says that John said something with the same import as the demonstrated utterance. Of course, since this utterance is itself difficult to interpret, it will not be easy for an audience to determine the truth or falsity of the claim about John. The claim itself, however, contains no ambiguity. If (2) is the unique logical form of (1) then the presence of ambiguity in (1) is mysterious (Rumfitt 1993: 435-6).³

Suppose, then, that SF is false. Can the defender of PT dispense with it and revise the theory in such a way as to require only LF? Ian Rumfitt has recently attempted to do just this (Rumfitt 1993). He begins by recasting the theory in the language of transformational grammar. Transformational theory supposes that the process of interpreting a sentence involves the construction of various mental representations of it, each codifying a different level of structure, and each derived from the others by transformational processes. D-structures represent a sentence's abstract phrase structure, S-structures correspond closely to its ordinary surface structure, and LF-structures make its quantificational form explicit. Rumfitt's proposal is that, while indirect speech is not paratactic at the level of D-structure and S-structure, it assumes a paratactic form at the level of LF-structure. For example, the English sentence 'Galileo said that the earth moves' has the S-structure

(3) [_{S₀}[_NGalileo][_{VP}[_Vsaid][_{S'}[_{COMP}that][_{S₁}[_{NP}the earth][_{VP}moves]]]]].

And Rumfitt's suggestion is that there is a transformation rule (he calls it *arboreal fission*) which maps this structure onto the LF-structure

(4) [_{S₀}[_NGalileo][_{VP}[_Vsamesaid][_{NP}that]] > [_{S₁}[_{NP}the earth][_{VP}moves]].

³ Michael Hand has recently drawn attention to two other cases which pose a very similar problem for SF. See Hand 1991.

The S_1 node becomes the head of a new free-standing tree, the complementizer is deleted and a demonstrative, referring to the new tree, is supplied to complete the VP of the original one. It is merely a *coincidence* that English uses the same orthographic form for both complementizer and demonstrative.

This proposal, Rumfitt claims, resolves the problems noted above. It concedes to Segal and Speas that SF is false, while maintaining that PT offers a correct account of the logical form of indirect discourse. Moreover, by detaching SF from LF, it permits the generalization of PT to languages, such as French, in which the surface form of indirect speech does not even *appear* to be paratactic. Rumfitt also claims that it solves the ambiguity problem. For if the second sentence of a paratactic pair is an LF-tree generated by arboreal fission, and if structural ambiguities in content sentences are resolved *before* fission, then it will be possible to construct distinct LF-representations of a structurally ambiguous report. For example, letting S_1 and S_2 stand for LF-representations corresponding respectively to the interpretations ‘Someone heard the hunters shooting’ and ‘Someone heard the hunters being shot’, we can read (1) as either⁴

(5) [John samesaid that]_S > S_1

or

(6) [John samesaid that]_S > S_2 .

The ambiguity in (1) is now explained as arising from the audience’s uncertainty as to which of these readings represents the speaker’s intentions (Rumfitt 1993: 437).

2.

Does arboreal fission save PT? There is an immediate problem, as Rumfitt recognizes. For, according to PT, the demonstratum of a paratactic demonstrative is an *utterance*. But according to Rumfitt’s revised version, it appears to be an abstract linguistic type – an LF-structure. And to suppose that the object of the indirect saying relation is a linguistic *type* is to revert to something like a quotational theory of indirect discourse. Rumfitt rejects such theories, of course, citing the usual reasons (most notably, failure to allow for the existence of context-dependent expressions in content clauses).⁵ To remedy the situation, he proposes the following ‘somewhat radical step’:

We have long been accustomed to the idea that tree-structures such as those labelled by ‘ S_1 ’ and ‘ S_2 ’ might *represent* disambiguations of an ambiguous string. What we need to introduce, I think, is the idea that we could also *use*

⁴ Using $[S]_S$ to signify a representation of the LF-structure of the English sentence S.

⁵ For an outline of a quotational theory, see Quine 1962: 212. For objections, see Davidson 1969, McDowell 1980.

such tree-structures to speak unambiguously ... we can imagine that the utterances in question are actually effected by inscribing (or otherwise ‘pronouncing’) the trees S_1 and S_2 . (1993: 437-8)

Suppose someone utters the sentence ‘John said that someone heard the shooting of the hunters’, and that we interpret this utterance as an assertion that someone heard the hunters shooting (the interpretation corresponding to LF tree-structure S_1). Then, Rumfitt claims, the utterance will be true ‘just if John stands in the two-place saying relation to the ensuing utterance of S_1 ’ (1993: 438).⁶

Now this is somewhat mysterious. It would be natural to say that an utterance u is an utterance of a linguistic type T just in case u consists in the production of a *token* of type T . So, to identify the referent of a paratactic demonstrative with an utterance of an LF-tree would be to identify it with the production of a token LF-tree (that is, with the act of inscribing a sentence in LF notation or of vocalizing a sentence in some phonetic version of this notation). This seems to be what Rumfitt has in mind when he talks of ‘inscribing’ or ‘pronouncing’ LF-trees. But this view looks very unattractive. When and by whom are the demonstrated tokens produced? It cannot be necessary for a speaker actually to *employ* LF notation in order to make successful indirect speech reports – otherwise virtually all such reports would be unsuccessful. But if reporters themselves do not produce the tokens, who does? And where are these tokens located? Of course, a *theorist* may produce token LF-trees in the course of specifying the logical form of an indirect speech report. But only in the very rarest of cases will it be to *these* tokens that the reporter was referring.

The only remotely plausible option seems to be to say that we are dealing here, not with public inscriptions, but with *mental representations*, and that the demonstrated ‘utterance’ is not in fact an overt speech act but a tokening of a mental representation. Suppose that meaningfully asserting the English sentence ‘Galileo said that the earth moves’ involves, among other things, tokening a mental representation of the LF-structure of the sentence ‘The earth moves’. And suppose I now assert that Galileo said that the earth moves. Then the claim will be that my utterance has the LF form

(7) [Galileo samesaid that]_S > m_{kf}

where m_{kf} denotes a tokening in my head of a representation of the LF-structure of the sentence ‘The earth moves’, and ‘samesaid’ is a two-place relation between a subject A and a token mental representation m such that ‘ A samesaid m ’ is true iff A produced an utterance u such that u has the same import as m .

This story may not be very attractive, and I doubt if it is the one Rumfitt has in mind, but does it work? In particular, does it circumvent the ambiguity trap? Its chances do not look good. Rumfitt, remember, holds that the ambiguity in an utterance such as (1) arises from uncertainty as to the referent of its LF-level

⁶ There is a slight terminological problem in Rumfitt's exposition, which I have silently corrected.

demonstrative. But suppose I utter (1), intending to assert that John said that someone heard the hunters shooting. Then, according to the present story, my utterance has the LF form

(8) [John samesaid that]_S > m_{kf1}

where m_{kf1} denotes a tokening in my head of a mental representation of the LF-structure S_1 . And here there is no uncertainty as to the referent of the demonstrative. For there is plainly only one suitable token mental representation available for demonstration – namely, the one tokened in my head in the course of my meaningfully uttering (1). So if my audience takes my utterance to refer to a token mental representation at all, then they will know which one is in question. Of course, there are many things about this token that they will not know. They will not know, for instance, which LF-structure type it is a token of. They will thus be unable to think of it under the same mode of presentation as I think of it. But it is generally agreed that it is not necessary for two speakers to think of an object under the same mode of presentation in order for them to understand each other's demonstrative references to it. (Compare (1) with a case where I secretly look at a playing card, place it face-down on the table, and then, pointing to it, say 'John chose that card'.) Here the paratactic theorist, by insisting that reference to content clauses is demonstrative in character, cuts the ground from under his own feet.

It might be objected that in the case described my audience would not in fact be able to identify the demonstrated mental token and so would not be able to establish the sort of informational access to it necessary for them to understand my report. But even if this were so, they would still not find my report *ambiguous*. For, in that case, rather than being uncertain which of two or more objects was the referent of my demonstrative, they would be unable to find *any* suitable referent for it. (If they could not locate m_{kf1} , then it is most unlikely that they would be able to locate any other candidate.⁷) They would accordingly be unable to assign any determinate interpretation to my report, which would be to them, not ambiguous, but incomprehensible.⁸

To sum up: Rumfitt claims that (1) is ambiguous because a hearer will be uncertain as to the identity of the referent of the LF-level demonstrative. But such uncertainty will be possible only in cases where there exists a plurality of candidate referents. And if the referent is an utterance, such cases will be very rare.⁹ Typically, a

⁷ The audience can be expected to know that speakers do not normally token multiple incompatible mental representations of their own utterances, and thus that only one candidate token representation is likely to exist.

⁸ This assumes a 'Russellian' view of singular reference. On a non-Russellian view, the lack of a referent would not render the report incomprehensible; but neither would it render it ambiguous.

⁹ Of course, we can imagine cases in which there exists a plurality of candidate utterances (at a party, say, where lots of people are talking simultaneously); but we do not want to say that it is only in such cases that an utterance of (1) would be ambiguous.

hearer will either be able to pick out the correct referent or be unable to pick out any. In neither case will they find the report ambiguous.

Now it may seem that I have been misrepresenting Rumfitt. For although he does describe the referent of a paratactic demonstrative as an utterance of an LF-tree, he does not ultimately endorse the idea that paratactic demonstratives refer to utterances at all. Instead, he follows McFetridge (1975) in detecting *deferred ostension* here, similar to that which occurs when we identify a type by demonstrating a token.¹⁰ McFetridge's suggestion (a response to the so-called 'counting problem') is that the 'that' of 'says that' ultimately refers, not to the immediately ensuing utterance, but to the set of utterances *bound* by that utterance (that is, to the set of utterances which can be indirectly reported by it). McFetridge calls such sets *propositions*.¹¹

This appears to offer an alternative explanation of the ambiguity in reports such as (1). A hearer might find such a report ambiguous, not because they were unable to identify the utterance which was the immediate referent of the LF-level demonstrative, but because they could not identify the *proposition* which was its ultimate, deferred, referent. The ambiguity might arise from uncertainty as to the course taken by the ostensive act after it had, so to speak, *passed through* the immediately indicated utterance.

This suggestion is, I think, on the right lines.¹² In its present form, however, it still fails to avoid the ambiguity trap. For a McFetridgean proposition can be identified only indirectly, by indicating an utterance which binds it. So in order for an audience to understand a demonstrative reference to a proposition, they will first have to identify the utterance which binds it. If they are unable to do this, they will find the reference incomprehensible.¹³ Moreover, each utterance *by definition* binds just one proposition (see Rumfitt 1993: 450). So if an audience can correctly identify the binding utterance, then they will be able to identify (at least by description) the unique proposition which it binds. In this case they will find the reference unambiguous. The only case in which an audience will remain uncertain as to the identity of demonstrated proposition is one in which there exists a variety of candidate binding utterances. And it was precisely the rarity of such cases that led us to reject the previous proposal.

Now, it may be objected that I have been working with too impoverished a notion of what an utterance is. An utterance, it might be said, is not just the production of a sentence token, but the production of a sentence token *with certain communicative intentions*. Suppose I utter (1), and that, as PT claims, the second term of the verb in this sentence refers to my utterance of the content clause. Now (the suggestion might

¹⁰ The term *deferred ostension* is Quine's, of course; see Quine 1969: 40ff.

¹¹ Rumfitt makes a slight modification to McFetridge's account; see note 16 below.

¹² If it is what Rumfitt has in mind, however, then it is odd that he should introduce the idea of *uttering* of LF-trees at all. For the point of that move was to show how the ambiguity in (1) could arise from uncertainty as to the identity of the utterance that is the *immediate* object of ostension.

¹³ Again, I assume a Russellian view of singular reference; see note 8 above.

go), in order for an audience properly to understand this reference, it will not be enough for them merely to identify the token action in question (to distinguish it, by its spatio-temporal location, say, from all other token actions). Rather, they will need to identify a certain canonical *description* of this token action – to decide whether to think of it as *the utterance of a sentence intended to be interpreted in accordance with LF-tree S₁*, or as *the utterance of a sentence intended to be interpreted in accordance with LF-tree S₂*. If they are uncertain which description is appropriate, they will find the reference – and consequently my report – ambiguous.

The problem with this suggestion, however, is to see how it can be incorporated into the paratactic theory. For PT holds that content clauses in indirect speech are *demonstrated*. And we cannot demonstrate items *under descriptions*. Provided that a hearer is able to identify a certain object, they are equipped to understand a demonstrative reference to it. That they cannot supply a canonical description of the object is irrelevant.

3.

The story so far is depressing. Rumfitt's attempt to hold on to the core paratactic thesis LF while avoiding the ambiguity trap has proved unsuccessful. The recurring problem has been that of finding a suitable referent for the putative demonstrative in an indirect speech report. Consider the kind of thing we want for a such a referent. If we are to avoid the ambiguity trap, then there must be a plurality of candidate referents for the demonstrative in sentences such as (1). So the referent had better not be a token utterance. Rather, we should look for deferred ostension to an abstract item of some kind. On the other hand, if PT is not to collapse into a quotational theory, the referent must not be a linguistic type, but must be contextually situated or indexed in some way. Finally, the referent cannot be a McFetridgean proposition; otherwise we fall again into the ambiguity trap.

Is there any other kind of abstract item that might fit the bill? I think there is. Consider the action type consisting of the act of producing tokens of a particular sentence type, S, in a particular context, C. It might, for example, be the act of my producing tokens of 'Today is my birthday' on a particular day. And suppose that C is such that every possible token of S produced in C will report every other.¹⁴ Call such an action type an *utterance type*.

Utterance types can be identified by means of pairs of sentence types and contexts. (Note that utterance types are not *identical* with such pairs; they are action types.) Not every pair of sentence type and context will pick out an utterance type, of course. For there will be many sentence-context pairs of which it is not true that all possible tokens of the sentence produced in the context would report each other. (Consider for example the pair of the sentence 'It is eleven o'clock' and the current

¹⁴ 'Report' here is a two-place relation between utterances, such that *u* reports *u'* just in case *u* could be used in a paratactic report to identify what was said by *u'*. The term is McFetridge's.

day.¹⁵) Now, I want to suggest that it is utterance types that are demonstrated in indirect speech reports.

Recall Rumfitt's story about arboreal fission. The idea was that at the level of S-structure the content clause of an indirect speech report is structurally integrated into the report, but that at LF-level it splits off. Consider again the sentence

(9) Galileo said that the earth moves.

At S-level, the form of this report is

(10) [_S₀[_NGalileo][_{VP}[_Vsaid][_{S'}[_{COMP}that][_S₁[_{NP}the earth][_{VP}moves]]]]].

But at LF-level, Rumfitt claims, it is transformed into

(11) [_S[_NGalileo][_{VP}[_Vsamesaid][_{NP}that]]].

(This is the complete LF-form of the *report*; the content clause has fissioned.) The core idea is that at some structural level reporting clauses and content clauses are doing very different things. And this has a sort of plausibility. After all, the reporter *asserts* the one but not the other (except when they assert them both of course – as in 'I assert that the earth moves'). Of course, assertoric force is not a structural feature, but we might reasonably expect the systematic difference in force between reporting and content clauses to correlate with some structural feature.

In addition, arboreal fission introduces a demonstrative into the reporting clause. Now, as we have seen, Rumfitt plays around with the idea that the referent of this demonstrative is itself something existing at the LF-level of structure – a type or token of an LF-tree generated from the content clause by the fission process. But there is no need to follow him here. We can suppose that the referent is just the speaker's utterance of the content clause – the act of producing the words 'the earth moves' in the context of the report. The fact that these words are part of a larger speech act which, at one level of structure, forms an indivisible whole, is irrelevant. The words are uttered and can be independently demonstrated.

So, there is nothing to stop us saying that

(12) Galileo said that the earth moves

has the LF form

(13) [Galileo samesaid that]_S > *u*

¹⁵ When a sentence-context pair defines an utterance type, I shall say that the context is *well-defined* for the sentence.

where u denotes the speaker's utterance of the words 'The earth moves'. But if the demonstratum is just u , then, of course, we fall straight into the ambiguity trap. So, let us suppose that act of demonstration passes, by deferred ostension, through u to the *utterance type* of which u is a token. This utterance type can be picked out by the pair of the sentence type of which a token is produced in the course of uttering u , and the context C in which u is uttered. So the logical form of (12) will be:

(14) [Galileo samesaid that]_S >U

where U denotes the utterance type picked out by the pair ⟨'The earth moves', C ⟩, and 'samesaid' is a two-place relation between a subject and an utterance type, such that 'A samesaid U ' is true iff A produced an utterance u such that u reports (has the same import as) an utterance of type U . It must be stressed that (14) does not say that Galileo *made* an utterance of type U , but that he *samesaid* an utterance of that type – that is, made an utterance reportable by an utterance of that type.

The ambiguity trap can now be circumvented in the following way. Suppose that the context which identifies an utterance type may include information about the intentions of the utterer. Then the act of uttering the sentence 'Someone heard the shooting of the hunters' with the intention that it should receive interpretation in accordance with LF-structure S_1 constitutes a different utterance type from that of uttering the same sentence with the intention that it should receive interpretation in accordance with LF-structure S_2 . Call these two utterance types U_1 and U_2 respectively. Since utterance types are abstract items, U_1 and U_2 will always be available for demonstration.

Now we can explain the ambiguity in (1) as arising from uncertainty as to the route taken by the act of deferred demonstration after it passes through the utterance of the content clause. A hearer who is uncertain whether it is U_1 or U_2 that is ultimately being demonstrated in (1) will find the report ambiguous. For they will be uncertain whether its logical form is

(15) [John samesaid that]_S > U_1

or

(16) [John samesaid that]_S > U_2 .

4.

That's the story, then. I think it is the only way to make PT run as a claim about logical form, while at the same time avoiding the ambiguity trap.¹⁶ Let us see how it

¹⁶ It is notable that Rumfitt eventually arrives, though by a rather different route, at a position significantly similar to the one advocated here. We noted that he endorses McFetridge's claim that

fits into the wider picture. It would be unwise to claim that the account offers an especially intuitive account of the logical form of indirect discourse. But if SF is false, then *no* version of PT is going to do that.

Does the proposal have anything going for it, beyond the fact that it saves a paratactic form for indirect discourse?¹⁷ Well, PT has its virtues and the proposal retains most of them. For instance, it preserves what Davidson calls *semantic innocence* (the view that embedding an expression in the content clause of an indirect speech report does not change its meaning). On the present proposal, this remains true, since after fission a content clause receives interpretation as a free-standing utterance.

Another virtue of PT was that it found semantic structure in the objects of verbs of indirect speech. This is important, of course, since by substituting for *p* in the schema 'A said that *p*' it is possible to produce indefinitely many novel, yet readily comprehensible, sentences. Unless the substituted expressions refer to items with semantic structure, our ability to understand such sentences will be mysterious. Does this virtue survive in the present revision? – that is, do utterance types have semantic structure? I don't see why not. We can suppose that they inherit the semantic structure of their corresponding sentence types.

Another attractive feature of PT was its ontological minimalism. It specified the truth conditions of indirect speech reports without mentioning senses, propositions or other intensional items. Our revised version is not quite so minimalist, referring as it does to action types. But these are harmless things. We can samesay the act of *uttering a sentence*, just as we can enjoy the act of *playing chess* and dislike that of *getting up*. It might be objected that the reference to contexts in specifying utterance types reintroduces a good deal of ontological baggage. For how can we specify contexts in a minimalist fashion – without, for example, invoking the apparatus of *possibilia*? There is a neat dodge we can employ here. We can think of a context as defined in terms of the reporting relations between utterances. (The reporting relation, remember, is defined in terms of the samesaying relation, which, following Davidson, we can treat as a semantic primitive.) We can think of a context for a sentence as defined by a set of axioms specifying disambiguating communicative intentions for the sentence, together with assignments to any context-dependent

the ultimate referent of a paratactic demonstrative is not an utterance, but the proposition bound by the immediately demonstrated utterance. Rumfitt points out that McFetridge's account fails to deal adequately with some counterfactual reports and suggests that we think of utterances as binding *propositional acts* instead. An utterance *u* is an instance of the propositional act *A* bound by another utterance *u'* iff *u'* either reports *u* or would report *u* if *u* existed. Rumfitt then introduces *propositions* as abstract items associated one-to-one with propositional acts. Finally he proposes that the 'that' of 'says that' refers to the proposition bound by the ensuing utterance (which, however, he still takes to be an utterance of an LF-tree; see his 1993: 447-49). Rumfitt's propositional acts have some similarities with what I call utterance types.

¹⁷ Saving a paratactic *form* for indirect speech is in itself a trivial exercise: as Blackburn once pointed out, we can just suppose that paratactic demonstratives refer to *propositions* (Blackburn 1975: 185).

expressions it contains. Then we can say that a set of axioms X describes a well-defined context C for a sentence S iff any two possible utterances of S would, when interpreted in accordance with X, reciprocally report each other.

Of course, a semantic theory's being minimalist isn't sufficient for its being true. And there are, in any case, other minimalist proposals on the market, which aim to achieve much the same end as PT without the bother of parataxis (see, for example, Higginbotham 1986; Larson and Segal 1995, ch.11). It is not clear, however, that these accounts are any less complicated than the present one, or that they are better equipped to avoid versions of the ambiguity trap. At any rate, I am content for the argument of this paper to take the form of a conditional. If we want to save PT as a claim about logical form, then *this* is how it must go. The reader is welcome to accept this conclusion as a *reductio* of PT.¹⁸

References

- Blackburn, S. 1975. The identity of propositions. In *Meaning, Reference and Necessity*, ed. S.Blackburn, 182-205. Cambridge: Cambridge University Press.
- Davidson, D. 1969. On saying that. In *Words and Objections: Essays on the Work of W.V.Quine*, ed. D.Davidson and J.Hintikka, 158-74. Dordrecht: Reidel. Reprinted in Davidson 1984, 93-108.
- Davidson, D. 1984. *Inquires into Truth and Interpretation*. Oxford: Oxford University Press.
- Hand, M. 1991. On saying that again. *Linguistics and Philosophy* 14: 349-365.
- Higginbotham, J. 1986. Linguistic theory and Davidson's program in semantics. In *Truth and Interpretation*, ed. E. LePore, 29-48. Oxford: Blackwell.
- Larson, R. and Segal, G. 1995. *Knowledge of Meaning*. Cambridge, Mass.: MIT Press.
- McDowell, J. 1980. Quotation and saying that. In *Reference, Truth and Reality*, ed. M. Platts, 206-37. London: Routledge.
- McFetridge, I.G. 1975. Propositions and Davidson's account of indirect discourse. *Proceedings of the Aristotelian Society* 76: 131-45.
- Quine, W.V.O. 1962. *Word and Object*. Cambridge, Mass.: MIT Press.
- Quine, W.V. 1969. *Ontological Relativity and Other Essays*. New York: Columbia University Press
- Rumfitt, I. 1993. Content and context: the paratactic theory revisited and revised. *Mind* 102: 429-454.
- Segal, G. and Speas, M. 1986. On saying $\delta\alpha\tau$. *Mind and Language* 1 (2): 124-132.

¹⁸ I am grateful to George Botterill, Jenny Saul and Peter Smith for comments on earlier drafts of this paper and to Gavin Boyce and Ian Rumfitt for discussions about the paratactic theory.