

QUOTATION: COMPOSITIONALITY AND INNOCENCE WITHOUT DEMONSTRATION

ANDREW BOTTERELL
Department of Philosophy
Sonoma State University
andrew.botterell@sonoma.edu

ROBERT J. STANTON
Department of Philosophy
University of Western Ontario
rstainto@uwo.ca

SUMMARY: We discuss two kinds of quotation, namely indirect quotation (e.g., ‘Anita said that Mexico is beautiful’) and pure quotation (e.g., ‘Mexico’ has six letters). With respect to each, we have both a negative and a positive plaint. The negative plaint is that the strict Davidsonian (1968, 1979a) treatment of indirect and pure quotation cannot be correct. The positive plaint is an alternative account of how quotation of these two sorts works.

KEY WORDS: propositional attitudes, semantics/pragmatics boundary

RESUMEN: Discutimos dos tipos de citas, a saber, citas indirectas (por ejemplo, “Anita dijo que México es bonito”) y citas puras (por ejemplo, “México” tiene seis letras). Hacemos dos planteamientos, uno positivo y otro negativo, con respecto a cada una. El negativo es que el tratamiento estrictamente davidsoniano (1968, 1979a) de las citas indirectas y puras no puede ser correcto. El positivo consiste en dar una explicación alternativa de cómo funcionan estos dos tipos de citas.

PALABRAS CLAVE: actitudes proposicionales, división entre semántica y pragmática

1. *Introduction*

In one sense, this paper has a narrow focus. We discuss two kinds of quotation, namely indirect quotation (e.g., ‘Anita said that Mexico is beautiful’) and pure quotation (e.g., ‘Mexico’ has six letters). With respect to each, we have both a negative and a positive plaint. The negative plaint is that the strict Davidsonian (1968, 1979a) treatment of indirect and pure quotation cannot be correct. The positive plaint is an alternative account of how quotation of these two sorts works; this account builds on some insights by various neo-Davidsonians, but it departs from Davidson’s specific view in key ways. (Even with respect to this narrow focus, we confess that we leave open numerous details for later exploration.) One might reasonably complain that this focus is *too* narrow. Happily, in another sense the paper is quite broad. It addresses foundational topics in philosophy

of language such as the nature of opacity and compositionality, the semantics/pragmatics boundary, and the difference between describing communicative activities and providing a semantics for specific linguistic devices. So, although the narrow issue of indirect discourse and pure quotation is in the foreground, these four larger issues are (often enough) in the background.

The paper is structured as follows. In this introduction, we will illustrate the kinds of cases we are interested in, and explain a question that they pose, viz., does quotation of these two sorts mean that semanticists can't have both compositionality and innocence? In the next section we present Donald Davidson's ingenious (negative) answer to that question, and we rehearse several problems for his account, some familiar, some novel. Finally, we present our own (negative) answer to the question, drawing on work by various neo-Davidsonians, and we consider a number of objections to that answer.¹

As noted, our focus in this paper will be on pure quotation and indirect quotation. These are illustrated in (1) and (2) below.

1. *Indirect quotation*

- (a) Winston Churchill said that Herr Hitler is a scourge on Europe.
- (b) Winston Churchill said that German Chancellor Adolf Hitler is a scourge on Europe.
- (c) Context of substitution: Winston Churchill said that ___ is a scourge on Europe.

2. *Pure quotation*

- (a) 'Twain' contains five letters.
- (b) 'Clemens' contains five letters.
- (c) Context of substitution: '___' contains five letters.²

¹Since readers could be misled by our choice of label, let us clarify what we mean by 'neo-Davidsonian'. We mean only that our account builds on insights from certain neo-Davidsonians such as Higginbotham, Larson, Ludlow, Segal, and especially Pietroski. As will emerge below, the spirit (as opposed to the pedigree) of our positive proposal is arguably closer to early Tarski, at least as he is understood by Gómez-Torrente (2001, pp. 145–146), than it is to Davidson (1968, 1979a).

²Some philosophers might deny that this can be a context of substitution. But they would be wrong to deny this. See Stainton 2000 for discussion.

In each case, the first sentence is, we suppose, true, while the second is false, despite substitution of co-referring singular terms.

These constructions are of special interest because they raise issues about opacity, compositionality and innocence. That is, a semantics for a language L can be viewed as a function that assigns semantic values to well-formed expressions of L. Two constraints are often placed on the project of providing a semantics for a language L. The first constraint, the constraint of *compositionality*, holds that the semantic value of an expression E of L is a function of E's constituent parts and the way those parts are put together. The second constraint, the constraint of *innocence*, holds that the semantic value of an (unequivocal) expression E of L does not vary depending on what context E is embedded in. As is well known, however, the existence of opacity seems to suggest that these constraints cannot be jointly satisfied. Consider the following two sentences:

3. Twain wrote *Huck Finn*.
4. Clemens wrote *Huck Finn*.

Working within a traditional truth-conditional semantics, and noting that 'Twain' and 'Clemens' are co-referential, each of 'Twain' and 'Clemens' will be assigned the same semantic value in the context of (3) and (4), respectively, namely the person Clemens/Twain. So, our semantics predicts that (3) and (4) have the same semantic value. So far, so good. Complications arise, however, when we introduce contexts that give rise to the phenomenon of opacity. By way of illustration, consider the following two sentences:

5. John said that Twain wrote *Huck Finn*.
6. John said that Clemens wrote *Huck Finn*.

Sentence (5) seems to have as a constituent part the name 'Twain', a genuine singular term; and (6) seems to have as a constituent part 'Clemens', also a genuine singular term. And, as we have said, these names are co-referential. So if the constraints of compositionality and innocence are true, it would seem to follow that (5) and (6) must themselves have the same semantic value. For if the constraint of compositionality is true, then the semantic values of (5) and (6) are functions of the semantic values of their constituent parts and the way those parts are put together. And if the constraint of innocence is true, then the constituents of (5) and (6) have the same semantic

values here as they do in (3) and (4) —i.e. the difference in parts makes no difference. Nonetheless, it is arguable that (5) and (6) do *not* have the same semantic value. For it certainly seems that John could say that Twain wrote *Huck Finn*, and yet fail to say that Clemens wrote *Huck Finn*.

This seems to suggest that there is something wrong with the semantic picture with which we began. In particular, this seemingly suggests that either the semantic values of some expressions are *not* functions of the semantic values of their constituent parts and the way those parts are put together, or that some (unequivocal) expressions can shift their semantic values. Now, the first disjunct seems out of the question: we are, after all, talking about the semantics of the types, and there really isn't anything else to the type beyond what its parts are, how they are put together, and what they mean. So, famously, Frege (1892) pursued the second, non-innocent, option. The question is: *does quotation, whether of the indirect quotation kind or of the pure quotation kind, really force the semanticist to give up at least one of compositionality or innocence?* (Call this “the big question”.)

2. Davidson's Negative Answer

Partly in response to this question, Donald Davidson (1968) proposed his ingenious paratactic account of indirect speech. (Later extended, in 1979a, to pure quotation.) According to Davidson, a sentence like (6) is really composed of *two* sentences, namely, ‘John said that’ and ‘Clemens wrote *Huck Finn*’. Furthermore, according to Davidson, the ‘that’ in the sentence ‘John said that’ functions as a demonstrative that picks out a token of the sentence ‘Clemens wrote *Huck Finn*’ that immediately follows it. How is this supposed to help? The problem, recall, is that our semantics predicts that (5) and (6) ought to be semantically equivalent, and this seems false. But Davidson's paratactic proposal resolves this problem. For if the logical forms of (5) and (6) are, respectively, (5') and (6') then the semantic values of (5) and (6) can diverge.

5'. John said that. Twain wrote *Huck Finn*.

6'. John said that. Clemens wrote *Huck Finn*.

For, as noted, on Davidson's view, ‘that’ in (5) demonstrates one object, namely a token of sentence (3), whereas ‘that’ in (6) demonstrates a distinct object, namely a token of sentence (4). And since

there is no incompatibility between John's being related to a token of sentence (3) and not to a token of sentence (4) —the two tokens being distinct objects— Davidson is at least in principle capable of making the right predictions about the truth-conditions of sentences (5) and (6).

As will be explained in detail below, according to Davidson 1979a (p. 90),³ pure quotation works essentially the same way. Quotation marks are referring expressions —specifically demonstratives. Such marks refer, to expressions and other shapes, by pointing out some utterance or inscription, in context, having the shape. Thus we get a negative answer to the big question, on both the indirect and the pure fronts.

2.1.1. Against the Paratactic Approach to Indirect Quotation

There are reasons to think that Davidson's paratactic analysis of indirect quotation will not do. The most serious has to do with Davidson's suggestion that 'that' is, or functions as, a demonstrative in sentences like (5) and (6), repeated below.

5. John said that Twain wrote *Huck Finn*.

6. John said that Clemens wrote *Huck Finn*.

Let's start with some grammatical facts. It appears that 'that' can be deleted in sentences like (5) and (6). Thus, (5*) has a reading on which it means the same as (5):

5*. John said Twain wrote *Huck Finn*.

Also, if 'that' is syntactically a demonstrative, occupying a noun phrase position, then it might seem that we should be able to replace 'that' with another demonstrative pronoun and still end up with a grammatical sentence. But neither (7) nor (8) is grammatical:

7. *John said this Twain wrote *Huck Finn*.

8. *John said it Twain wrote *Huck Finn*.

Finally, the demonstrative 'that' has different phonological properties from the (*sometimes* homophonous) complementizer 'that'. Specifically, as Speas and Segal (1986) point out, only the complementizer can be phonologically reduced to *th't*. Thus, although (9) is well-formed, (10) is not:

³ Page numbers are from Davidson 1984.

9. John said *th't* Marie is cute.
10. *Marie is cute. John said *th't*.

What this suggests is that ‘that’ is *not* grammatically a noun phrase in (5) or (6). In which case, it’s not the demonstrative noun phrase [NP that]. What appears in (5) is not [NP that] but [COMP that], which can be deleted, can be phonologically reduced, and cannot be replaced by [NP this] or [NP it]. That is, ‘that’ in discourse reports seems to function syntactically as a clause- or sentence-introducer. Thus, ‘that’ belongs with such expressions as the ‘whether’ and ‘if’ of indirect questions, as well as the gerundive ‘...’s ...ing’ of, for example, ‘I regret Smith’s leaving’. Syntactically, all of these are complementizers, not noun phrases. What appears in (11), in contrast, is [NP that], which cannot be deleted, cannot be phonologically reduced, and can be replaced by [NP this] and [NP it].

11. Twain wrote *Huck Finn*. John said that.

This shows that one cannot argue *from* the grammar of ‘that’ to Davidson’s semantic conclusion. One cannot, for instance, argue as follows: “Because the same word appears at the end of (11) and in the middle of (5), it’s reasonable to conclude that demonstration is involved in both cases.” For, given the above considerations, it seems quite *unlikely* that the same word appears in (11) and (5). Of course, at bottom all Davidson is committed to is the claim that ‘that’ functions *semantically* as a demonstrative in sentences like (5) and (6). And maybe he can support this conclusion even if ‘that’ in such sentences is not *syntactically* a demonstrative. As against this, however, consider the following examples—in which ‘that’ in indirect discourse reports doesn’t seem to function like a demonstrative, even semantically. The first example is due to Higginbotham (1986). Suppose one of us writes (12) on the blackboard, and while pointing to it, says (13):

12. He is a nice fellow.
13. Every boy believes that.

This activity cannot be used to convey that every boy has a good opinion of himself. And yet this is clearly something that can be done by saying (14):

14. Every boy believes that he is a nice fellow.

What this suggests is that ‘that’ in (14) does not function as a demonstrative, even semantically. If it did, it should always be understood the way (13) is: as saying, of some contextually identified particular boy, that every boy believes him to be a nice fellow.

Here is a different example, due to Michael Hand (1991). It’s clear that, *pace* Davidson’s proposal, (15) and (16) do not mean the same thing. (Indeed, it’s not clear that (16) means anything at all.)

15. I didn’t say that there was any beer in the fridge.

16. I didn’t say that. There was any beer in the fridge.

In a similar vein, consider that in VP ellipsis constructions, the referent of the explicit demonstrative in the trigger-sentence is carried over to the gapping site. An example will clarify what we mean. (17) cannot mean that Utpal saw some contextually salient thing, ‘that’, but Luis didn’t see a *different* contextually salient thing.

17. Utpal saw that, but Luis didn’t.

Instead, (17) can only mean that Utpal saw some particular thing, and Luis didn’t see that same thing. (That is, (17) does not have all the readings of ‘Utpal saw that, but Luis didn’t see that’, since the ‘that’s in this latter conjunctive sentence *can* refer to distinct things.) Keeping this in mind, consider sentence (18):

18. Silvia said that she was hungry, but Lucia didn’t.

This sentence *can* mean that Silvia said that Silvia was hungry, but Lucia didn’t say that Lucia was hungry. Yet, on Davidson’s view, (18) has the following logical form:

19. Silvia said that, but Lucia didn’t. She was hungry.

From (17), we know that, since ‘that’ is supposedly a *demonstrative* in the “trigger-sentence” of (19) (i.e. ‘Silvia said that’), its reference must be carried over to the gapped site (i.e. ‘Lucia didn’t ___’). So what Lucia didn’t say must be the very thing that Silvia *did* say: the explicit ‘that’ in the trigger-sentence must refer to the same thing as the “covert” demonstrative in the ellipsis-site. Hence it is wrongly predicted that this sentence can only mean that Lucia didn’t say that *Silvia* was hungry.

What is common to all three example-classes is this: there is some kind of semantic dependence between the main clause (i.e. the one

which precedes and contains the ‘that’, and describes the agent and the saying-relation), and the subordinate clause (i.e. the one which captures the thing-said). In Higginbotham’s case, this dependence is binding by a higher quantifier of a lower variable; in Hand’s case, it is licensing of ‘any’, in the subordinate clause, by an explicit negative particle in the main clause; and in the final case, it is a trigger-sentence licensing a gap-site in the subordinate clause. In these cases, and numerous others that one could construct, the semantic dependence is not well captured by the mechanism of across-sentence demonstration.

So, not only are there problems with the claim that ‘that’ is grammatically a demonstrative in sentences like (5) and (6), there are also problems with the weaker claim that ‘that’ is merely functioning semantically as a demonstrative in such sentences.

It is worth noting one other problem with Davidson’s proposal, since a similar issue will dominate when we discuss pure quotation. Davidson’s suggestion is that ‘that’ in sentences like (5) and (6) really is a *demonstrative*. The thing about demonstratives is that their referent is whatever thing is made appropriately salient in the environment. Thus ‘That is ugly’ can be used to talk about a salient phone, an indicated car, a dog being pointed at, or what-have-you. What ‘that’ picks out depends upon what object is made salient. In which case, one would expect the word ‘that’ in indirect discourse reports similarly to have a flexible, highly context-sensitive, reference to something “sayable”. But this just is not the case. For instance, (20) cannot be used to assert that Alice said that Bob Dole is a great patriot —adding that, in fact, Dole is a goof:

20. Alice said that Bob Dole is a goof.

Even if the speaker utters (20) while pointing at a token of ‘Dole is a great patriot’, written on a blackboard, (20) cannot be used to assert that Alice said Dole is a great patriot. In contrast, (21) certainly can be used to make just this complex of statements:

21. Bob Dole is a goof. Alice said that. [Speaker points at inscription of ‘Dole is a great patriot’.]

(The example is taken from Stainton 1999 (p. 269), which defends it against obvious countermoves.) In indirect speech reports, then, the word ‘that’ cannot be used to “demonstrate” anything other than the subordinate clause that immediately follows —in ways very

unlike the use of a true demonstrative [_{NP} that]. What this suggests is that there isn't really pragmatically determined ostension going on at all. Instead, there is an inflexible semantic rule of the kind we will propose below.

2.2. Against the Paratactic Approach to Pure Quotation

Before criticizing it, let us rehearse some of the details of Davidson's story about pure quotation by illustrating it with an example. According to him, the quotation marks in (22) refer to the expression set-off in (23); and they do so by pointing to a token of this shape—namely, the three-lettered token right after the open-quote.

22. 'Jam' has three letters.

23. Jam

The predicate 'has three letters' then applies to the reference of the quotation marks. Since this object—namely, the expression 'jam'—satisfies this complex predicate, sentence (22) as a whole comes out true. Crucially, then, this demonstrative theory of quotation does indeed preserve both innocence and compositionality, and it permits a negative answer to our big question.

We turn now to criticisms of Davidson's views on pure quotation. In doing so, we should note that others have complained about having quotation marks *refer*, having them refer to *shapes*, and having the demonstrata be *linguistic tokens* (i.e. inscriptions or utterances). These are all good points. Our concern, though, lies elsewhere: it has to do with the *means by which* reference is purportedly achieved. Specifically, we deny that quotation marks function like demonstratives.

What exactly is the problem? Curiously enough, three consequences of the demonstrative theory, which Davidson himself notes, leads directly to it. Davidson writes:

A. "Any token may serve as target for the arrows of quotation. . ." (1979a, pp. 91–92)

B. "the question of location [of the quoted material] is trivial". (1979a, p. 90)

C. "Quotation is a device for pointing to inscriptions (or utterances) and can be used, and often is, for pointing to inscriptions or utterances spatially or temporally outside the quoting sentence." (1979a, p. 91)

That is, if quotation marks are demonstratives, they can be used to demonstrate *any* contextually salient token. As a result, the location of the demonstrated token is immaterial; in particular, the demonstrated item may fall outside the quoting sentence. Theses (A) through (C) are, we agree, implied by the demonstrative account. The problem is, none of them is true *about quotation understood as a language-internal device*. In particular, the location of the quoted material *is not* immaterial: quotation marks cannot demonstrate just any contextually salient token. From which we conclude, *contra* Davidson, that quotation marks are not demonstratives.

To illustrate our contention that location matters, consider two examples. Example one: it is, we take it, a datum that (24) does *not* have a reading on which it is true —no matter what the context.

24. ‘Undoubtedly’ starts with the 19th letter of the alphabet... strangely enough.

What the quotation marks in (24) refer to can only be the shape ‘undoubtedly’; and this shape begins with the 21st letter of the alphabet, not the 19th. Were Davidson right, however, (24) *would* be readable as true: if, as (A)–(B) would have it, quotation marks really can refer to *any* salient token, then the quotation marks in (24) can, in appropriate contextual circumstances, refer to the token of ‘strangely enough’. Which would yield a reading of (24) such that it says roughly the same thing as (25):

25. Strangely enough. The expression of which that is a token undoubtedly starts with the 19th letter of the alphabet.

So read, (24) would come out true.

What this first example highlights is that, in quotation, location matters. Taking the demonstrative theory seriously, the issue of *how* the demonstratum is contextually determined immediately arises; and with it, the possibility that the demonstratum might well be something other than the material within the marks. But, we maintain, this is not a real possibility —which means that quotation marks do not function like demonstratives. Here’s another way of making the same point. If quotation marks really were demonstratives, capable of referring to *any* salient token, the following should be well formed, and readable as true:

26. ‘___’ has nine letters... obviously

(García-Carpintero 1994 and Gómez-Torrente 2001 both independently noted this latter feature of Davidson's view.)

The problem, in short, is that Davidson treats quotation marks as garden variety demonstratives. But if quotation marks were garden variety demonstratives, then it seems that context could determine some sentence outside the marks as the demonstratum. And that simply is not possible. Put otherwise, and to anticipate a point that will loom large below, Davidson seems to assimilate quotation to "talk about talk" generally; but, we insist, quotation constitutes a special device, a special convention, for speaking about words. And as such, it is subject to special constraints that ordinary "talk about talk" is not. This is what Davidson's view doesn't capture. (Note: what Davidson says in (A)–(C) is likely true, for the most part, of communication about words. But pure quotation of the kind found in (2) and (22) must be distinguished from the—very general—ability to "talk about talk".)

2.3. Objections to Our Criticism of Davidson, and Replies

We end this negative section by considering three plausible replies on behalf of Davidson's original proposal. (We will focus on pure quotation, but the replies could be used, *mutatis mutandis*, to defend Davidson's paratactic theory for indirect speech as well.)

First reply. The reason (24) is not read as true has merely to do with pragmatics. In which case, there is no problem with the idea that quotation marks are demonstratives—recalling, of course, that they refer to the shape of *the most salient* utterance/inscription. The thing is, this reply suggests that, if only the context were right (in particular, if only 'strangely enough' *really were* the most salient token), then (24) would be read as (25) is. But that just isn't so: there just aren't any contextual devices, like pointing, that can override the salience of the words between the quotes. (Note too that, for reasons made familiar by Grice, hearers try quite hard to interpret speech such that the thought communicated is true; and, if Davidson's view were correct, such a true thought should be available. Still, it is never seized upon by readers.)

Second reply on Davidson's behalf. Could it not be that quotes are demonstratives, but that there is a salience convention governing their use? For example, a convention to the effect that the most salient token *is always the one inside the quotes*? To say this would be to recognize that, in some sense, location is more than pragmatically relevant. (Hence Davidson's (A) through (C) would not be

strictly correct for pure quotation, however insightful they are about talk-about-language in general.) But it would not amount to giving up the idea that quotes are demonstratives. To which we say: there are two ways of understanding ‘convention of salience’ here. On the one hand, one may take the purported convention to be merely heuristic, like the usual rules for assigning reference-in-context. (Compare: ‘she’ refers to the unique perceptually or conversationally salient woman in the context. This rule is defeasible. Sometimes ‘she’ refers to a man, sometimes to a boat, etc.; and sometimes the referent is not antecedently salient in the context; e.g., ‘She certainly is taking her time’, said of a perfectly non-salient waitress, to break a conversational silence.) On the other hand, one may assume the convention to be inviolable. Going the first way, the heuristic should be capable of being overridden. In which case, (24) and (26) should have readings on which they are true—which they don’t. Pursuing the second option, however—i.e. saying that the reference of the quotes is *invariably and necessarily* the material inside them—two worries present themselves. First, Davidson is the last person to believe in inviolable conventions for reference fixing, so a tension is introduced with his larger philosophy of language. (See his 1979b, 1982 and 1986.) Second, introducing such a convention actually amounts to assigning reference by something quite unlike pointing to tokens. After all, a demonstrative whose reference is no more pragmatically elastic than ‘Mexico City’ is no demonstrative at all. Put otherwise, one central attraction of Davidson’s idea was that it explained away an apparent “reference shift” (i.e. within quotation) in terms of something understood independently: a demonstrative taking on a different reference, depending on which contextually salient particular is demonstrated. On the present proposal, this is now lost. It seems, then, that the demonstrative theory cannot be saved by appeal to a salience convention: to give such a rule is, we believe, to *abandon* the demonstrative theory of quotation; it is, indeed, to adopt something more like our own proposal, given below.

Here is a third possible reply. It builds on the second. Mightn’t quotation marks be like ‘I’ and ‘now’, indexicals whose in-context referent *cannot* vary depending on speaker’s intentions, pointings, etc.? This would explain why quotation marks always pick out the material inside them; and the problem of “rogue utterances” appearing on blackboards or elsewhere in the sentence doesn’t arise after all since, on this proposal, quotation marks are *pragmatically inflexible* indexicals. The problem with this proposal is that we can specify

the complete proposition expressed by —and in consequence, the context-independent truth value of— the expression type in (22):

22. ‘Jam’ has three letters.

This observation holds for quotation-containing sentences generally. However, sentences containing pure indexicals like ‘I’ and ‘now’ do not have context-independent truth-values; to the contrary, they aren’t true or false except given a context of utterance.⁴ This suggests that quotation marks do not function like pure indexicals. Indeed, this observation gives rise to a more general worry for Davidson since on his view, *no quotational sentence type should express a context-independent truth-condition, because they all contain token-reflexives*. Because he applies parataxis to all such cases, they should all be like sentences containing indexicals. But they clearly are not.

Perhaps it will be replied that all this shows is that quotation marks ought to function more like anaphors or bound-variable pronouns, with the rule being that the quotation marks are somehow *bound* to the material inside the quotation marks. The suggestion would thus be that quotation marks work in a manner similar to ‘itself’ in (27), yielding a sentence that *is* true or false independent of a context of utterance:

27. Every number is divisible by itself.

The problem with this proposal, however, is that it essentially concedes the point at issue rather than refuting it. For anaphors and bound-variable pronouns are *semantically* linked to their antecedents by a context-invariant compositional rule; they are not pragmatically linked. To treat quotation marks on the model of anaphors or bound-variable pronouns, then, is to grant our point that the link between quotation marks and the material inside them is not pragmatic. While not strictly a notational variant of on our own semantics-oriented functional solution (to be introduced below), this proposal abandons the spirit of the Davidsonian pragmatic story for quotation in much the way our proposal does.

⁴Granted, there are some sentence types that contain indexicals *whose every spoken token* is true, such as ‘I am speaking now’. But despite this, the sentence type itself lacks a truth value, and is not about any individual or any time.

2.4. Background Issues, Part One: Semantics/Pragmatics and Linguistic Devices

One of our central questions in this paper is whether quotation, of either the indirect or the pure variety, forces the semanticist to abandon at least one of compositionality or innocence. As we've seen, Davidson gives a negative answer, but his particular approach faces a raft of problems. With respect to indirect quotation, the grammar of the 'that' of 'says that' does not support treating this word as a demonstrative noun phrase; various dependencies between the matrix clause and the subordinate clause after 'that' militate against treating it as (not syntactically but still) semantically demonstrative; and the referent of the 'that' of 'says that' doesn't seem to be fixed by salience. With respect to pure quotation, when there is something inside the quotation marks, they cannot refer to some "rogue utterance" outside; the marks cannot appear without something between them; and there are standing sentences containing quotation marks.

Our overall negative conclusion is thus that quotation marks do not function like demonstratives. Ditto for [_{COMP} that]. This negative result relates to two of the background foundational issues flagged at the outset, and it's worth bringing them out explicitly here. Davidson, maybe unwittingly, essentially treats quotation as a *pragmatic* phenomenon, in the sense of involving demonstration of a salient entity using a demonstrative expression. He thus implicitly places quotation on the pragmatic side of the semantics/pragmatics boundary: it belongs with context-based reference assignment, not with type-based compositional semantics. This strikes us as an important mistake. 'That' of 'believes that' and quotation marks are conventional, properly linguistic devices that function without the need of extra-linguistic contextual clues.⁵

Similarly, Davidson (1979a) overtly assimilates pure quotation to "talk about talk". But, as hinted above, there is a world of difference between describing the communicative activity of making claims about words, and providing the semantics of a particular language-

⁵ Curiously, Davidson is not alone in being subject to this kind of complaint. Any account which has the link between item-quoted and quotation-as-a-whole being a matter of pragmatic demonstration will face similar worries. For example, Reimer 1996—who is a neo-Fregean about quotation, and who gives up on innocence—is nevertheless in the same boat with Davidson in this regard. She has quotation marks serving only a clarificatory role, like a pointing gesture, and this seems to us equally wrong.

specific device that allows one to do this. We can highlight the contrast we have in mind by means of a couple of comparisons. First, there are lots of ways to convey notions of quantity: with gestures, adverbs of quantification, pictures, etc. Providing a formal semantics for quantifier words like ‘each’ and ‘no’, words that appear in the determiner position of phrases, will not account for all these other ways of getting across claims about quantity. But that’s a good thing. A formal semantic theory of quantifying determiners would miss its mark, qua semantic theory about natural language, if it *did* apply to quantitative communication in general. Or again, on a more clearly related note, think of all the ways of getting across what some person Jaime believes. One can say, ‘For Jaime, p ’ or ‘As far as Jaime is concerned, p ’. In answer to ‘Does anyone really think that p ?’, one could point at Jaime. One could even perform a charade to convey what Jaime thought about p . Any of these actions would do the job. Granting this, if someone missed the difference between a compositional semantics for the construction ‘ S believes that p ’ and all these other ways of conveying what someone believes, we would straighten them out. Our position is the same for ‘said that’ and quotation marks: they are special devices, a semantics for which should not accommodate every imaginable means of talking about language.

This contrast is important for another reason. It helps us avoid objections that have been raised to previous accounts. Washington (1992) in particular notes that we can talk about expressions by employing strategies like off-setting on the page, intonation, or just context. All of these can make it clear that an expression is being displayed or presented, rather than being used in the normal way. Washington then complains that views which focus on quotation marks as contentful devices miss these cases. (See also Reimer 1996.) What he notes is correct. In particular, *our* semantics for quotation marks, given below, will not apply to all these cases. But, contra his complaint, that is perfectly in order —because it’s far from obvious that Washington has presented us with a *linguistic* natural kind. A certain communicative activity, “talk about talk”, is occurring in all such cases; but, to repeat, quotation marks and ‘said that’ are specialized devices. More than that, they operate in a highly constrained, convention-based way, resulting in a complex entity that refers via a recursive mechanism. So, a theory of the latter device need not, and *should* not, extend to every example of the former practice.

3. *Our Neo-Davidsonian Alternative*


Davidson's paratactic proposal has two parts. The first is that 'that' and quotation marks contribute not the semantic values of parts, but rather certain linguistic objects. The second part is that 'that' and quotation marks function as demonstratives. For reasons that emerged above, we want to abandon the idea that quotation involves demonstration. But we want to retain the idea that the relata are individuated, at least in part, by syntactic features. Moreover, we want to do this without introducing reference shifting for the items in question. What is required is a mechanism that allows this. We will here propose just such a mechanism. We begin with indirect discourse, and then turn to pure quotation.

3.1. Compositionality and Innocence *without* Demonstration: Indirect Quotation

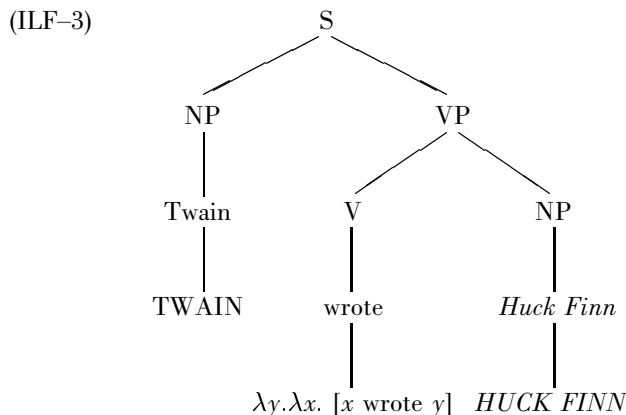
The idea, in a nutshell, is that 'that' is a complementizer, not a demonstrative noun phrase; and its contribution is determined by a semantic rule, not by ostensive demonstration; nonetheless, the Complementizer Phrase refers to *syntactic* material contained in the whole sentence. (For closely related proposals, which consider many more details than discussed here, see our neo-Davidsonian predecessors: Higginbotham 1986, Larson and Ludlow 1993, Pietroski 1996, and Stanton 1999.) The way this works, we think, is that $[_{COMP} \text{that}]$ denotes a function. The input to this function is (in part)⁶ the syntactic material that is its sister. The output, which is denoted by the whole Complementizer Phrase rather than by $[_{COMP} \text{that}]$ on its own, is the Interpreted Logical Form (hereafter 'ILF') associated with the syntactic structure of $[_{COMP} \text{that}]$'s sister node. (The same story can be told when the COMP node is empty, as in $[_S \text{Jaime said } [_{CP} [_C e]] [_S \text{it was raining}]]$.)

What is an ILF? Let us represent the *un*interpreted logical form of sentence (3) using a tree structure, as follows:

⁶This proviso is important, for another input to the function will have to be contextually salient objects for any indexicals that occur in the complementizer phrase. Thus, for instance, in 'John said that she is married', the syntactic structure 'she is married' can't be the only input to the function denoted by 'that'; the salient woman needs to be input as well. For the moment, it suffices to note that these considerations suggest that the function denoted by 'that' is not *merely* a function from syntax. We leave it open what other sorts of things the function might take as arguments. We are indebted to Paul Pietroski for raising this issue.



The *interpreted* logical form of sentence (3) is obtained by pairing each terminal node of the uninterpreted logical form with its ordinary semantic value. This gives us:



Now, suppose that what agents are claimed to be related to in sentences like (5) and (6) are neither sentence tokens, nor the semantic values of sentences, but ILFs. On this view, since the ILF associated with (5) is distinct from the ILF associated with (6) —the former containing $[_{NP}$ Twain] while the latter contains $[_{NP}$ Clemens]— a believer could be related to one ILF and yet not be related to the other ILF. So, (5) could be true while (6) is false —which is precisely the result we want. For instance, consider again sentence (5):

5. John said that Twain wrote *Huck Finn*.

Taking a leaf from Higginbotham, our proposal is that (5) will be true just in case John uttered an object suitably similar to the ILF of (3):

3. Twain wrote *Huck Finn*.

Thus, our proposal is that in a sentence of the form $\ulcorner X \text{ says that } P \urcorner$ the semantic value of $[\text{COMP that}]$ will be a function that takes the uninterpreted logical form of ' P ' as argument and yields the interpreted logical form of ' P ' as value.

3.2. Compositionality and Innocence *without* Demonstration: Pure Quotation

Let us turn in detail to the phenomenon of pure quotation. We begin by rehearsing the by-now familiar problem posed by quotation, applied this time to pure quotation. Here it is:

28. *The Puzzle of Pure Quotation*: The reference of an expression of the form ' α ' appears not to depend solely upon the reference of the α in question.

(Using corner quotes, rather than talking explicitly about expressions of a certain form, this would be: $\ulcorner \alpha \urcorner$ appears not to depend solely upon the reference of $\ulcorner \alpha \urcorner$.) Evidence for this comes from sentences with the following form:

29. ' α ' has three letters

Put (30) in, and the resulting sentence is true; put (31) in, however, and the sentence produced is false.

30. Jam

31. Preserves

The problem is, despite the fact that (30) and (31) are evidently co-referential, (32) and (33) do not refer to the same thing.

32. 'Jam'

33. 'Preserves'

In which case, the reference of expressions of the form ‘ α ’ must depend on something more than the reference of the α in question.

As noted near the outset, one solution to this puzzle of opacity, made famous by Frege, is to allow the reference of words to shift according to context. Applied to this case: sometimes (30) refers to sweet sticky compote, and sometimes it refers to a word —namely, ‘jam’. Which thing (30) refers to depends upon how it is embedded. Specifically, and most importantly for present purposes, in a context like (29), (30) shifts its reference from jam to ‘jam’. Of course a similar point holds for (31): sometimes it refers to preserves, and sometimes it refers to ‘preserves’. How does this explain opacity? Well, on this proposal, (30) and (31) actually are *not* co-referential —at least not in quotational contexts. In such a context, the first refers to ‘jam’, but the second does not. So the puzzle dissolves: the reference of expressions of the form ‘ α ’ really *does* depend exclusively upon the reference of the α in question —it’s just that what α itself refers to is inconstant. Crucially, however, once the reference of the α in question is determined, so is the reference of the expression of the form ‘ α ’. (See Washington 1992 and Reimer 1996 for sophisticated contemporary variants.)

There is, of course, an awkward consequence of this solution to (28): it sacrifices *innocence*, i.e. the aforementioned idea that a(n) (un-equivocal) word always has the same semantic value, no matter what the context. We consider this too high a price. Taking a leaf from certain neo-Davidsonians, we will now propose our own innocence-preserving solution to the puzzle of pure quotational opacity. Once again we will appeal to a semantic rather than a pragmatic mechanism, and to a functional expression rather than a referring term.

Let us begin with an obvious point: it is not a problem *per se* if the reference of a whole expression does not depend exclusively upon the reference of its parts. To take but one example, no one is concerned by the fact that, despite their parts being co-referential, (34) and (35) fail of co-reference.

34. The half-brother of John’s mother.

35. The mother of John’s half-brother.

These examples are unproblematic because what compositionality demands is merely that whole-reference be a function of part-references *together with syntax*. And (34) and (35) are, of course, syntactically different. Returning to the case of quotation: compositionality does

not require that the reference of things of the form ‘ α ’ should be determined by the reference of α alone; it dictates instead only that the reference of ‘ α ’ should depend upon the reference *and syntax* of the whole. Thus (28), it seems to us, isn’t a really a puzzle at all—at least not as it stands. What *would* be a real puzzle is (36):

36. *The Revised Puzzle of Pure Quotation*: The reference of expressions of the form ‘ α ’ appears not to depend solely upon the reference of its parts, even together with its syntax.

This would be a real puzzle because it seems that this is all language users have to go on. So if meaning depended on more than this, it would be a mystery how speakers and hearers figure out what whole expressions mean.

But, so far as we can see, (36) isn’t even *prima facie* plausible—because what accounts for the difference in reference of, for example, (32) and (33) is precisely their distinct syntax. (A related point about compositionality appears in Pietroski 1999, p. 248.) It remains to explain *how* differences in syntax make for differences in the reference of things having the form ‘ α ’. Here is our proposal. We take from Davidson the focus on quotation marks, rather than on the expressions inside the marks. That is, we do not have the latter shifting reference in context, but look instead to the quotation marks, understood as more than mere contextual markers, to explain why the subject of the sentence is an expression, and not the thing usually denoted by the expression. Here we build on the idea introduced in the previous section: instead of having quotation marks *refer to expressions*, we propose to treat them as denoting *functions*. Specifically, quotation marks denote a function from a symbol to itself (i.e. the identity function, with domain restricted to symbolic items).⁷ Crucially, however, the input to the quotation function is *the symbol inside the marks*. The input is *not* the denotation of the symbol. Because it receives the symbol between the marks as input, the output of the function will be that very symbol. This output symbol then serves as the argument for the function denoted by the

⁷ After completing this paper, we had our attention drawn to Parsons (1982), which seemingly first introduced the idea that quotation marks might stand for the identity function. In addition, that paper addresses the relationship of this idea to Frege’s own views, and lays out some alternative implementations which we have ignored here. The details of Parson’s view are, not surprisingly, rather different from ours, and he does not develop or defend the identity function idea to nearly the same extent. But one of the central ideas is clearly present.

relevant meta-linguistic predicate (e.g., ‘has three letters’), yielding the correct truth conditions.

Working through an example should help clarify the proposal. Take (22), repeated below:

22. ‘Jam’ has three letters.

To preserve innocence, the denotation of ‘jam’ must invariably be a certain sweet sticky food. To get the right truth conditions, however, the *word* ‘jam’ must be what is contributed, rather than jam itself. The issue is, how is this achieved? How does this node manage to contribute the word itself to the function denoted by the meta-linguistic predicate ‘has three letters’, without ‘jam’ shifting its reference? Our proposal is that the identity function denoted by the quotation marks pays attention to the *form*, rather than the denotation, of the node to its immediate right. Specifically, it takes the syntactic item in (30) as input, and yields that very item as output. Which gives the desired semantic value for (32). You might nonetheless ask: how does our proposal distinguish (32) and (33), repeated below, given the co-reference of (30) and (31)?

30. Jam

31. Preserves

32. ‘Jam’

33. ‘Preserves’

Well, the reference of each of (32) and (33) will depend upon the semantic value of the quotation marks (i.e. an identity function whose domain contains only symbols) and the *syntax* of (30) and (31) respectively. But, obviously, these two are syntactically distinct. The complexes (32) and (33) therefore contribute different entities to the function $[h: (\lambda x)h(x)$ is T of x iff x has three letters], because the *input* to the function denoted by the marks is distinct.

We want to stress two things about this proposal. First, the input is determined by concatenation/embedding, syntactic relations *par excellence* —hence there’s no violation of compositionality in the only sense that is empirically well motivated. (We’ll have much more to say about this immediately below.) Second, the symbol inside the marks does (typically) have a denotation —indeed, it has the same

denotation it always has.⁸ Hence innocence is preserved. It's just that the denotation is semantically otiose, since it is not the argument for the function. True enough, the sense in which the reference of expressions of the form ' α ' depends on syntax is a bit different from the sense in which the reference of (34)–(35) does. In the latter case, syntax amounts to word order; whereas in the case of quotation, “paying attention to syntax” amounts to having the linguistic form, rather than its denotation, serve as input to a function. But, we insist, there is no violation of compositionality if the reference of ' α ' can be made appropriately sensitive to syntax in either sense. And, treating quotation marks as functional expressions whose input is the form of the concatenated symbol, this is easily done. Nor, to repeat, is there reference shifting: for example, (30) everywhere—in every linguistic context—refers to jam. And the quotation marks have a constant reference as well: to the identity function. So, compositionality is retained, *and innocence as well*. Thus we too can give a negative answer to the big question.

3.3. Objections to Our Positive Proposal, and Replies

3.3.1. A Complaint about ILFs

Some readers will have qualms about Interpreted Logical Forms, and their adequacy for representing, in indirect discourse, the thing which is said. It is worth noting, in reply, that our ILF proposal needn't be correct in every detail in order for the alternative framework that we are presenting to remain attractive. What is important for our purposes, both with respect to pure quotation and indirect discourse, is that:

- i) 'that' and quotation marks are functional expressions rather than referring terms;
- ii) the input to the function denoted by 'that' and quotation marks is determined by a semantic rule, not by pragmatic considerations; and
- iii) the output of the function can, when necessary, show sensitivity to the form of the input to the function, thereby yielding hyperintensionality.

⁸What happens when the item inside the marks lacks a denotation is an interesting question, which we put aside here.

Each of i–iii could be true even if some of the details of our ILF view are mistaken. For example, we admit with respect to the indirect discourse sub-case that the argument and the value of the function denoted by ‘that’ will have to include more than just interpreted and uninterpreted logical forms. In addition, phonological and orthographical information will be relevant to the determination of the precise value that the function denoted by ‘that’ yields. (This even holds for pure quotation: as García-Carpintero 1994 points out, ‘*jam*’ may exhibit an importantly different reference from ‘jam’.) What is important from our point of view is articulating a framework for understanding quotation and propositional attitude reports that retains some neo-Davidson’s insights—that form matters for hyper-intensionality, and that innocence is worth preserving—but rejects Davidson’s idea that the mechanism for this rests on reference of [_{COMP} that] or quotation marks via demonstration.

And this is why we prefer our Davidson-*inspired* account, according to which quotation marks denote a function whose input is invariably determined, independent of speech context, by the syntactic form of the expression inside the marks. This variant on Davidson’s original idea retains its advantages: i.e. explaining away the puzzle of opacity, while preserving innocence. But, unlike its precursor, our adaptation allows for the pragmatic inflexibility of quotation.⁹

3.3.2. The Whole Language Objection

The current proposal, with respect to indirect quotation, is that the function denoted by [_{COMP} that] applies to the embedded clause that follows it: it takes the uninterpreted logical form of that clause as argument, and yields as value the interpreted logical form of that clause. But, it might be objected, that function is nothing less than *the meaning function for the entire language*. More precisely, the

⁹ In the context of Davidson’s overall views, the move away from the demonstrative theory forces a question: are Davidson’s larger philosophical projects adequately served by treating quotation marks (and the ‘that’ of ‘says that’) as merely *similar* to demonstratives? For instance, would it be enough if these words denoted syntax-sensitive functions, rather than *extra-linguistic context* sensitive demonstratives? Or do his overall purposes dictate that they really *be* demonstratives (and hence *not* functional expressions)? To take but one example: we wonder whether the central role that Davidson assigns to tokens sits well with a “they’re only similar” answer; or again, we’re unsure whether one gets a reduction of quotational, discursive and propositional attitude opacity to contextual ostension if the relevant constructions don’t *really* contain demonstratives. Given the importance of this question, we really must try to answer it. . . some other time.

function denoted by ‘that’ must somehow have recourse to the meaning function that applies to the embedded clause that follows it. For how can the function denoted by ‘that’ take an uninterpreted logical form and yield an interpreted logical form unless it has access to the semantic theory for the entire language in question? But, continues the objection, it is implausible to suppose that a speaker who is competent in using the ‘that’ of indirect discourse is capable of grasping that function.

In response to this objection, however, let us simply note that everybody —from Davidson to Frege to Montague— who adopts a relational view of indirect quotation is saddled with the consequence that *some* mechanism makes available the meaning of the embedded clause following the ‘that’ of indirect discourse to users of the language in question. In Davidson’s case, for example, *the truth theory for the entire language* computes the meaning of the token utterance that follows ‘that’, and then that utterance is referred to. So, Davidson too is committed to the idea that speakers who understand the ‘that’ of indirect discourse must be capable of grasping the truth theory for the entire language in question. And this is certainly no better than our view. Our reaction to this objection, then, is to bite the bullet —denying that it’s implausible that ‘that’ *has recourse to* a compositional mechanism capable of computing the ILF of its complement. (To be clear, we do not identify the meaning of [_{COMP} that] with such a global compositional mechanism.)

Of course, someone might object that the problem is much more serious than ‘that’ having access to a rich compositional semantics. They might argue as follows: you are assuming that the meaning function M, employed by a user of L, yields meanings for all sentences, including sentences of the form ‘X says that P’. Given the principle of compositionality, however, the value of M for such an argument depends upon the meanings of the parts of sentences of that form, including the meaning of ‘that’, which, by hypothesis, encapsulates the meaning function M itself. Thus, continues the objection, there is ineliminable circularity and self-reference at the heart of the proposal.

In response, let us distinguish the phenomenon of self-reference from that of recursion. Our proposal certainly appeals to semantic rules that are recursive; but this is not to say that those semantic rules are circular or self-referential. For example, it’s not as if the ‘that’ in [_{CP} that snow is white] somehow already has stored up the meaning of the whole complex of which it forms a part —i.e.

that in computing the meaning of ‘that snow is white’ ‘that’ must already have access to the meaning of ‘that snow is white’. What ‘that’ has access to is the generative power to compute the meaning of its mother node, ‘that snow is white’, given the (form and) meaning of *its complement* ‘snow is white’. This is not self-reference. (Compare: in ‘it’s not the case that it’s not the case that snow is white’, the outermost ‘it’s not the case that’ has access to the meaning of ‘it’s not the case that snow is white’, which requires that the inner ‘it’s not the case that’ have access to the meaning of ‘snow is white’. But this does not give rise to a problem of circularity. The meaning of ‘snow is white’ is built up recursively from the meaning of ‘snow’, ‘is’ and ‘white’; that meaning is then used to yield the meaning of ‘it’s not the case that snow is white’; and that meaning is then used to yield the meaning of ‘it’s not the case that it’s not the case that snow is white.’) Thus, while on our proposal ‘that’ is always computing the meaning of the embedded clause that immediately follows it, this is not to say that the part-meaning makes illicit reference to the meaning of the whole of which it is a part.

3.3.3. The Iterated Embedding Objection

Speaking of recursion, one might object as follows. If quotation marks denote a symbol-to-symbol function—in particular, if quotation marks denote the identity function—then the output of the function should itself be able to serve, recursively, as the input to the quotation-function; and, because the function is an *identity* function, such a repeated application of the function should induce no change in the final output. But, continues the objection, quotation does not admit of *empty* recursion of this kind. So, the proposed account cannot be correct.

The point of the objection can be brought home with an example. It’s a datum that (37) does not have the same reference as (38).

37. ‘Jam’

38. ‘‘Jam’’

In particular, while the latter denotes something that begins with a quotation mark, what the former denotes begins with a letter. But, runs the objection, the identity-function account of quotation will assign equivalent denotations to (37) and (38). After all, (37) is equally well represented as (39); while (38) is captured by (40)—where, in each case, f is the identity function with domain restricted to symbols:

39. $f(\text{jam})$

40. $f(f(\text{jam}))$

But, of course, because f is the *identity* function, (39) and (40) are co-referential. Hence, this account erroneously predicts that (37) and (38), like their paraphrases (39) and (40), will co-designate. Generalizing: the identity-function proposal incorrectly makes iterated quotation look vacuous.

To understand where the objection goes wrong, recall that on our view the quotation function does not take as input the *denotation* of the symbol inside the marks. Instead, it takes the quoted symbol itself as input. Thus, in the case of (38), what the quotation function takes as input is the symbol in (37). Hence what it outputs is *this very symbol*, quotation marks and all. Put in terms of (39) and (40), the point is that these are *not* adequate representations of (37) and (38) respectively. The reason is, in (39) and (40) one naturally takes the denotation of the expression inside the braces as the input to the function. And, because f and $f.f$ are extensionally equivalent, (39) and (40) therefore stand for the same entity. It then seems that (37) and (38) must also refer to the same thing. But this functional notation misrepresents quotation because, to repeat, the input to the quotation function is determined by the form, and not the content, of the symbol within the marks. So, the objection from vacuous recursion fails. And the identity-function account of quotation stands.

3.4. Background Issues, Part Two: Compositionality and the Nature of Opacity

There are many notions of compositionality at play in linguistics and philosophy of language. For example, some understand compositionality to require that whole meanings always be determined as follows: given two sisters nodes, one of the two will stand for a function, the other will stand for an argument, and the function in question must apply to the semantic value of the node which contributes the argument to yield the meaning of the mother node. Another related understanding of compositionality takes it that the meaning of a whole must be exhaustively determined by what the parts are, and the *shape* of the tree that the parts belong to. It should be clear that the mechanism we suggest for quotation is not compositional in either of these senses. Our mechanism has the syntactic item being contributed, not just the denotation of the two nodes; and it isn't

just the shape of the tree that fixes whole meaning, but rather the items at the terminal nodes.

The existence of these and other restrictive notions of compositionality might cause one to worry that our view ends up violating compositionality, even though the whole point is to preserve both it and innocence. We think this worry is misplaced. It's misplaced because such highly restrictive notions of compositionality are without motivation. When we reflect a little on why we want natural languages to be compositional, it emerges that there are two related reasons. First, the semantic facts about a whole expression type must be fixed by the context-invariant rules of the language, plus context-invariant features of that expression. (Since we are talking about types, what else could fix content?) Given the nature of sentences, their features seem to be restricted to the form and content of the parts, how the parts are combined, plus certain global formal features of the whole (e.g., intonation patterns, mood, etc.) Thus the rules must operate on those features, to give the meaning of the whole. A second related reason for endorsing compositionality goes like this. The semantic facts about a whole expression must be such that we humans can compute those facts on the basis of things *we* know. This rules out algorithms that are infinite, or are unlearnable by us, etc. Hence, we want natural languages to be compositional because we want to account for *our* semantic competence. Each of these strike us good reasons for wanting compositionality of some sort—but they patently aren't good reasons for imposing compositionality of the highly restrictive sorts noted above. There are more facts about an expression type than what nodes are sisters, and what each denotes; there are, in particular, facts about an expression's tree that go beyond its shape. What's more, we humans know such additional facts. Specifically, we know what items sit at the terminal nodes. Hence there is no reason not to allow our semantic rules to pay attention to what forms occupy the terminal nodes, taking these as inputs for functions. This, of course, is just what we do.

One might worry that allowing functions to have access to the syntactic structure, including terminal nodes, of items in the syntactic tree will make compositionality trivial. Our first response is this. Given that compositionality is non-trivial when the semantics is allowed to pay attention to tree structure, syntactic categories of the parts, and content of the terminal nodes, it's hard to see how it becomes wholly trivial if the compositional rules can also pay attention to the form of the terminal nodes, and what this entails for the tree as a whole. To see this, consider a thought experiment.

Suppose you do not speak Spanish, but are told the tree structure, syntactic categories, content of terminal nodes, *and form-of-parts* for the sentence ‘Ya estuve allí y sé bien que no hay’. You still won’t be able to figure out what the whole sentence means, because you don’t know enough about how whole-meanings are compositionally determined in Spanish. What you lack are precisely the non-trivial compositional semantic rules.

Additionally, that compositionality remains non-trivial is shown by the fact that other “hard cases” for compositionality remain just as hard given our revised notion. Consider two familiar examples. Chomsky (1977) noted that the truth conditions we assign to ‘Mountains are climbed by fools like me’ and ‘Poems are written by fools like me’ are interestingly distinct. We hear the second as universal: all poems are written by fools. But we don’t hear the first this way. And yet the structures seem to be the same, and the part meanings suitably similar. Or again, John Searle (1978, 1980) notes truth conditional differences between things like ‘John cut the grass’, ‘John cut the cake’, and ‘John cut his finger’ that don’t seem to trace to structure or part meaning. Finally, what a structure contributes to meaning can vary widely from case to case; in particular, the meaning of the nominal-modifier complex frequently reflects things other than part meanings and structure. For instance, compare ‘Christmas cookie’ (“made to be consumed at”), ‘Girl guide cookie’ (“sold by”), ‘oatmeal cookie’ (“made of”), ‘yellow cookie’ (“coloured”), ‘fortune cookie’ (“containing”), ‘doggie cookie’ (“made to be eaten by”), and ‘Walmart cookie’ (“sold at”). (See also, among many others, Carston 2002, Jackendoff 2002, Pustejovsky 1995, Sperber and Wilson 1986, and Travis 1985 for related puzzles.) Such differences in intuitive truth conditions, without appropriate difference in structure or part-meanings, have recently lead some authors —e.g., Fodor (2001), Moravcsik (1998), Pietroski (2003)— to conclude that natural language semantics isn’t compositional after all. (Nor, of course, are they the first to draw this conclusion.) Tellingly for our purposes, these sorts of puzzles don’t simply *evaporate* as soon as we allow compositional rules also to have access to the form of the items at the terminal nodes. If these afford counterexamples to compositionality traditionally construed, they afford counterexamples to our narrower notion too. So the latter can’t be “trivial”.

So much for compositionality. When we reflect upon what motivates it, we see that our proposal is compositional *in any sense that human limitations warrant*. (That our story is not compositional in every imaginable sense is simply not of interest to us, working as we

do on the nature of human language.) Another important background point is about opacity. We hope it's clear how to extend the present account to the variety of opacity we find in direct discourse (e.g., 'John said 'jam''). There too, the quotation marks will stand for the identity function, and the input will be (in part) the syntactic item between the quote-marks. Given this, we can get an idea of what quotation in general, of all three types, has in common. Namely:

The Feature Cluster

- a) There is a function whose input is (in part) a syntactic item;
- b) which syntactic item becomes the input depends upon syntactic relations like sisterhood/concatenation, embedding, head-complement, etc. (e.g., appearing between certain marks, or being the grammatical complement of a complementizer);
- c) the function outputs something which is either itself a purely syntactic item (pure quotation and direct quotation), or is instead an interpreted syntactic item (indirect quotation).

We conjecture that something similar happens in several cases of genuine opacity, including especially propositional attitude contexts.¹⁰ Showing that must wait for another paper; still, in this light, let us briefly say something about our hidden motivations. On the face of it, while it is clear enough why it might be attractive to treat quotation marks as denoting a function from syntactic items to syntactic items, it might not be clear why 'that' should be similarly viewed as denoting a function from syntactic items to ILFs. Our motivations stem from two sources. For one, we find ILF accounts of propositional attitude reports attractive in general. The attractiveness of such accounts derives from their ability to show how it is in principle possible to reconcile *semantically* the phenomenon of opacity with the constraints of compositionality and innocence. The other source of motivation, however, is more speculative. If we are right, there isn't a quotational kind of opacity on the one hand and a propositional attitude kind of opacity on the other; rather, both kinds of opacity

¹⁰ Why 'several cases' and not 'all'? Because there are hard cases of conventionalized linguistic devices that don't obviously fit this mould. With regard to propositional attitudes, there are devices like 'As far as John is concerned...' With regard to pure quotation, there is the use of italics (*Jam* has three letters) and also "mixed quotation" ('John said that life is "hard to understand"'). Until we sort out how these might be assimilated to the "feature cluster", we prefer to stick to the weaker claim that several cases exhibit it.

involve ineliminable reference to form and structure. Moreover, if we are right then the constraints of compositionality (in the sense in which it is empirically motivated) and innocence can be satisfied in both cases without appeal to pragmatic considerations. If we are right, that is, then this *unitary* phenomenon of opacity counts as a genuinely *semantic* phenomenon.¹¹

REFERENCES

- Carston, Robyn, 2002, *Thoughts and Utterances*, Blackwell, Oxford.
- Chomsky, Noam, 1977, *Essays on Form and Interpretation*, North Holland, New York.
- Davidson, Donald, 1986, “A Nice Derangement of Epitaphs”, in LePore 1986, pp. 433–446.
- , 1984, *Inquiries into Truth and Interpretation*, Oxford University Press, Oxford.
- , 1982, “Communication and Convention”, in M. Dascal (ed.), *Dialogue: an Interdisciplinary Approach*, Amsterdam, John Benjamins. (Reprinted in Davidson 1984, pp. 265–280.)
- , 1979a, “Quotation”, *Theory and Decision*, vol. 11, pp. 27–40. (Reprinted in Davidson 1984, pp. 79–92.)
- , 1979b, “Moods and Performances”, in A. Margalit (ed.), *Meaning and Use*, Dordrecht, Reidel. (Reprinted in Davidson 1984, pp. 109–121.)
- , 1968, “On Saying that”, *Synthese*, vol. 19, pp. 130–146. (Reprinted in Davidson 1984, pp. 93–108.)
- Fodor, Jerry A., 2001, “Language, Thought and Compositionality”, *Mind and Language*, vol. 16, pp. 1–15.
- García-Carpintero, Manuel, 1994, “Ostensive Signs”, *Journal of Philosophy*, vol. 91, pp. 253–264.
- Gómez-Torrente, Mario, 2001, “Quotation Revisited”, *Philosophical Studies*, vol. 102, pp. 123–153.
- Hand, Michael, 1991, “On Saying That Again”, *Linguistics and Philosophy*, vol. 14, pp. 349–365.
- Higginbotham, James T., 1986, “Linguistic Theory and Davidson’s Program in Semantics”, in LePore 1986, pp. 29–48.

¹¹ Many thanks to Ernie Lepore, who discussed quotation at length with us, and to Ray Elugardo, who provided extensive comments on a previous draft. Thanks also to Dan Blair and Catherine Wearing for comments on the penultimate draft. This paper owes a substantial debt to Paul Pietroski’s writings on innocence and opaque contexts. We see it as a development and defense of his basic insights. See especially his 1996, 1999. (Though, roughly speaking, Pietroski has the complementizer *referring* to the ILF, rather than having it stand for a function, as we would.) Pietroski’s work, in turn, builds on Higginbotham 1986, Segal 1989, Larson and Segal 1995 and Larson and Ludlow 1993, among others. We gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada. Finally, a special acknowledgment to the two anonymous referees of *Crítica* for their helpful comments on a previous version.

- Jackendoff, Ray, 2002, *Foundations of Language*, Oxford University Press, Oxford.
- Larson, Richard and Peter Ludlow, 1993, "Interpreted Logical Forms", *Synthese*, vol. 95, pp. 305–355.
- Larson, Richard and Gabriel Segal, 1995, *Knowledge of Meaning*, The MIT Press, Cambridge.
- LePore, Ernest (ed.), 1986, *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, Basil Blackwell, Oxford.
- Moravcsik, Julius M., 1998, *Meaning, Creativity, and the Partial Inscrutability of the Human Mind*, CSLI, Stanford.
- Murasugi, K. and R. Stainton (eds.), 1999, *Philosophy and Linguistics*, Westview Press, Boulder.
- Parsons, Terence, 1982, "What Do Quotation Marks Name? Frege's Theories of Quotations and That-Clauses", *Philosophical Studies*, vol. 42, pp. 315–328.
- Pietroski, Paul M., 2003, "The Character of Natural Language Semantics", in A. Barber (ed.), *The Epistemology of Language*, Oxford University Press, Oxford, pp. 217–256.
- , 1999, "Compositional Quotation (without Parataxis)", in Murasugi and Stainton 1999, pp. 245–258.
- , 1996, "Fregean Innocence", *Mind and Language*, vol. 11, no. 4, pp. 338–370.
- Pustejovsky, James, 1995, *The Generative Lexicon*, The MIT Press, Cambridge.
- Reimer, Marga, 1996, "Quotation Marks: Demonstratives or Demonstrations?", *Analysis*, vol. 56, no. 3, pp. 131–141.
- Searle, John, 1980, "The Background of Meaning", in J. Searle, F. Kiefer and M. Bierwisch (eds.), *Speech Act Theory and Pragmatics*, Reidel, Dordrecht, pp. 221–232.
- , 1978, "Literal Meaning", *Erkenntnis*, vol. 13, pp. 207–224.
- Segal, Gabriel, 1989, "A Preference for Sense and Reference", *Journal of Philosophy*, vol. 86, pp. 73–89.
- Speas, M. and G. Segal, 1986, "On Saying δat ", *Mind and Language*, vol. 1, pp. 124–132.
- Sperber, Dan and Deirdre Wilson, 1986, *Relevance*, Harvard University Press, Cambridge.
- Stainton, Robert J., 2000, "Objects and Senses and Substitutions", *Dialectic*, vol. 39, no. 3, pp. 593–599.
- , 1999, "Remarks on the Syntax and Semantics of Mixed Quotation", in Murasugi and Stainton 1999, pp. 259–278.
- Travis, Charles, 1985, "On What Is Strictly Speaking True", *Canadian Journal of Philosophy*, vol. 15, pp. 187–229.
- Washington, Corey, 1992, "The Identity Theory of Quotation", *Journal of Philosophy*, vol. 89, pp. 582–605.

Received: September 29, 2004; revised: March 9, 2005; accepted: April 6, 2005.