

Quotation, Context Sensitivity, Signs and Expressions

Herman Cappelen and Ernie Lepore

Draft of June 8, 2006

Can one and the same quotation be used on different occasions to quote distinct objects? The view that it can is taken for granted throughout the literature (e.g. Goddard & Routley 1966, Christensen 1967, Davidson 1979, Goldstein 1984, Jorgensen *et al* 1984, Atlas 1989, Clark & Gerrig 1990, Washington 1992, García-Carpintero 1994, 2004, 2005, Reimer 1996, Saka 1998, Wertheimer 1999). Garcia-Carpintero (1994, p. 261) illustrates with the quotation expression "gone". He says it can be used to quote any of the following items:¹

1. The expression ("gone' is dissyllabic');
2. Different types instantiated by the tokens ("*gone*' is cursive');
3. Different types somehow related to the token (say, the graphic version of the uttered quoted material, or the spoken version of the inscribed quoted material, as in "gone' sounds nice');
4. Different tokens somehow related to the quoted token ('What was the part of the title of the movie which, by falling down, caused the killing? - 'gone' was');
5. The quoted token itself ('At least one of these words is heavier than 'gone" which you should imagine written in big wooden letters).

Here's another of his examples:

6. There are contexts in which the quotations '*Madrid*' and 'Madrid' would have the same content, but there are easily conceivable contexts in which they would have different contents... (García-Carpintero 1994, p. 260; cf. also García-Carpintero 2005, p. 97)

We can summarize his points (as well as those of the other authors referenced above) as the thesis of Quotation Context Sensitivity (QCS):

¹ Our numbering

QCS: Let S be a sentence with a quotation expression Q.² Two utterances u and u' of S can express different propositions because Q in u and in u' quotes different items.

In short: which item we use a quotation expression to quote can vary across contexts of utterance. The question we want to explore here is whether semantics or pragmatics can accommodate all quotational flexibility; put bluntly, whenever the u and u' of QCS quote different items must it always be on either semantic or pragmatic grounds? We will ultimately defend a negative answer to this question for an interesting subclass of data.

In what follows, we discuss several aspects of quotation incompatible with a semantic explanation of the QCS phenomena but compatible with a pragmatic one. However, we come to the conclusion that some variability data posited as instances of QCS resist even a pragmatic explanation. That result leaves an explanatory gap that we will fill – surprisingly – with a syntactic interpretation of the relevant puzzling data. This does not mean that we will be defending a *syntactic* interpretation of QCS, since to do so would presume that the distinct utterances we have in mind are of a *single sentence with a single quotation expression*. We want to deny this. Instead, we will argue that a proper understanding of the relevant data requires that we acknowledge something that has hitherto been overlooked – namely, that distinct sentences with distinct quotation expressions have been conflated in the service of QCS. Along the way, we'll endorse an obvious but also almost universally neglected minimal semantics for quotation, and a story about the scope and limits of our linguistic practice of quotation.

The paper divides into two parts. In Part I, we advance four considerations against a semantic construal of QCS. In Part II, we outline three pragmatic strategies for accommodating quotational variability data in general, and then present cases that seem to resist a pragmatic treatment; finally, we present a syntactic solution to the recalcitrant data.

Part I: Against a Semantic Construal of QCS

§1 Four concerns with a semantic QCS

One obvious strategy for explaining the variability data about quotational usage is to posit semantic context sensitivity (or ambiguity) for quotation expressions. In this

² without any context sensitive expressions other than possibly Q

regard, a quotational expression behaves semantically like normal indexical or demonstrative expressions, for example, like 'I', 'now' and 'that'. Two utterances of the same quotational sentence can disagree in truth value because the semantic value of its quotation expression changes across contexts of use. With that said we turn directly to problems confronting this suggestion.

§1.1 QCS and Disquotation

Our first worry about a semantic QCS is that it blocks an explanation of the most intriguing (and the most ignored) aspect of quotation: namely, its disquotational nature, as exemplified in (7):

7. "Quine" quotes 'Quine'.

More generally: *disquoting a quotation expression thereby retrieves whatever item it quotes*. Ordinary speakers, solely by virtue of being competent with the practice of quotation, know that disquotational quotation sentences are true.³ But any view compatible with a semantic QCS cannot guarantee their truth. Here's why: Call the first quotation expression Q1 in (7) and its second Q2. If quotations are semantically context sensitive, on any given occasion, all different kinds of items can be quoted by Q1 and Q2 respectively. Nothing therefore guarantees that whatever Q1 quotes on any given occasion will be in any way whatsoever related to whatever Q2 quotes on that occasion. Were quotation (semantically) context sensitive (as a semantic construal of QCS predicts), there would be innumerable ways to make (7) false. But to be told that we're not guaranteed the truth of a quotation sentence like (7), as a matter of meaning alone, will come as a surprise to most competent speakers.⁴

Here's a more straightforward way of getting at this point: the practice of quotation underwrites a Strong Disquotational Schema (SDS):

³ Indexicals can of course form sentence utterances of which are knowable a priori, such as 'I'm here now'. The a priori status of utterances of such sentences is inferred from the character of the contained indexicals. The proposals we are now considering with respect to the context sensitivity of quotation do not ascribe characters that would secure this kind of apriority to (7) or SDS. To see that, note that the theories we are considering would allow both 'Madrid' = 'Madrid' and '"Madrid" = "Madrid"' to be false. That said, we obviously don't mean to rule out the *in principle* possibility that context sensitive terms have characters that underwrite apriority.

⁴ The weight you should put on this intuition will of course depend, in large part, on your view of context sensitivity and how transparent you should expect the phenomenon to be for speakers. We are among those who are highly skeptical of views that postulate surprising context sensitivity (Cappelen/Lepore 2004, pp. 112-113) – the claim that there can be false utterances of (7) is one of those.

(SDS): “‘e’” alone quotes 'e'.⁵ (where ‘e’ stands for any quotable item whatsoever).⁵ It’s impossible to quote 'Quine' with, for example, "Jason"; only “‘Quine’” quotes ‘Quine’. Of course, you could name or describe or demonstrate a quotable item in countless ways; you can name ‘Quine’ with “‘Jason’” if the urge should move you.⁶ Since a semantic QCS is compatible with denying SDS, so much the worse for it.

§1.2 Proximity

A particularly close relationship stands between a quotation and its semantic value, one like no other kind of expression bears to its semantic values. Compare the relationship between ‘Quine’ and its semantic value, namely, 'Quine', with the relationship between the latter and its semantic value, i.e. the great American philosopher of the 20th century. The former are intimately related in a way that the latter are not. The quotation expression (i.e. "Quine"), in a way to be explained in Part II, has its semantic value, i.e. the item it quotes ('Quine'), contained in it, but the semantic value of 'Quine', i.e. the philosopher, is nowhere to be found in 'Quine'. Were it there, confrontations with tokens of 'Quine' would be a lot more interesting than they are.

Other authors couch this point in different ways. Here are three from the literature:

- i. ...a quotation somehow pictures what it is about (Davidson 1979 p. 82).
- ii. A quotation is...a hieroglyph. ..[that] designates its object...by picturing it (Quine 1940, p. 26).
- iii. ...we can go from knowing the quotation of any expression to knowing the expression itself (Saka 1998, p. 116).

⁵ SDS is the simplest and most natural account of quotation. It is endorsed by some authors in passing. DQR: For any expression e, the left quote (lq) followed by e followed by the right quote (rq) denotes e (Richard 1986, p. 397)

Richard uses 'denotes' for 'quotes' and talks of expressions rather than quotable items (we'll see the significance below), but the idea behind SDS is in DQR. Wallace (1972) commits himself to SDS:
 ..the denotation of the result of enclosing anything in quotes is the thing itself (Wallace 1972, p. 237)

SDS is also in Ludwig and Ray 1998, p.163, note 43; Mates 1972, p. 21; Salmon 1986, p. 6; Smullyan 1957; Gomez-Torrente 2005, p.129.

⁶ From which it follows that whatever quotations are they can't be – contrary to what an extensive literature on quotation has tried to convince us – names, descriptions or demonstratives.

We prefer to speak vaguely (at this stage) of a quotation expression *containing* whatever item it quotes. A theory of quotation must explain (or explain away) our intuitions, however vague or imprecise, about the special nature of this relationship. With a semantic explanation of QCS, it would be hard if not impossible to see how this explanation might go.

Proximity is obviously closely related to SDS: "Jason" can't quote 'Quine' *because* the two do not stand in the relevant proximal relationship to one another (regardless of whether this is spelled out via (i), (ii), (iii) or with our preferred idea of containment). If what a quotation expression quotes can vary from context to context, then what's quoted on any given occasion can't be included or contained in that quotation itself; and likewise, if quotation were semantically context sensitive, there could be no backwards road from a quotable item to a quotation. And so on.

In sum: in senses yet to be made clear, 'Quine' *is contained in* "Quine", "Quine" *does picture* 'Quine', and we *can go from knowing* 'Quine' to knowing that which quotes it and vice versa. But if quotation were semantically context sensitive, it would be impossible to explain any of these familiar features of quotation.

§1.3 Over-Generation

A third consideration against a semantic explanation of QCS is that it makes false predictions about what we can say using quotations, as least on a liberal construal of the flexibility data, as suggested by some of García-Carpintero's data (in particular, cases (4)-(5) above). On the assumption that quotation is context sensitive, it follows that what a quotation quotes on any given occasion is determined either by speaker intentions or what is contextually salient. There isn't anything else on offer. According to García-Carpintero (and others), quotation expressions quote not only types but also tokens, as in (4)-(5). We have no complaint against an audience being directed to a token with an utterance of a quotation expression. But, just the same, we find it quite difficult to retrieve true utterances of a quotation-sentence like (8):

8. 'a' ≠ 'a'

Were García-Carpintero right, it should be no harder to render an utterance of (8) true than one of (9):

9. That ≠ that.

If a speaker intends to refer to a token with the first utterance of "a" and to an expression-type with her second, then her utterance of (8) is true. Or suppose that prior to her utterance the speaker had been drawing our attention to differences between two tokens of the same quotation-expression. Again, the result should be a false utterance of (8). However, as a matter of fact, intuitively, we and our informants find it very difficult to get either of these readings; even after explicitly rendering both salient. Similar considerations extend to readings of other quotation-sentences whose truth seems to require quoted-tokens, as in (10)-(11):

10. 'I' didn't exist twenty seconds ago.

11. 'I' tastes like peach.

The impossibility (or at least great difficulty) we and our informants confront in retrieving true readings of utterances of (10)-(11) provides at least some evidence against an unrestricted semantic construal of QCS. Quotational usage, regardless of whether it's semantic or pragmatic, lacks the flexibility that an unconstrained semantic QCS predicts. However, not too much should hang on such reactions. These skeptical intuitions alone, even if correct, do not rule out a semantic explanation of some instances of QCS. They don't rule out, for example, the possibility that distinct types can be quoted on distinct occasions of use (and so, they don't rule out García-Carpintero's cases (1)-(3) and (6)). The next round of considerations is more forceful in this regard.

§1.4 Indirect Reports and Collectivity

We turn to our final objections to a semantic explanation of the data in support of QCS. In Cappelen/Lepore (2004), we pressed into service a number of tests in order to identify (semantic) context sensitivity. We'll discuss two (others are in Cappelen/Lepore (2004, Ch. 7)): the Inter-Contextual Disquotational Indirect Reporting and the Collectivity Tests.

'Inter-Contextual Disquotational Indirect Reporting' is just an ugly term for the following unexceptional phenomenon: Take an utterance *u* of a sentence *S* by speaker *A* in context *C*. An inter-contextual disquotational indirect report of *u* is an utterance *u'* in a context *C'* (where $C' \neq C$) of 'A said that *S*.' In Cappelen/Lepore (2004), we defended using such reports to test for context sensitivity.

Inter-contextual Disquotational Indirect Reporting Test: If an occurrence of an expression *e* in a sentence tends to block disquotational indirect reports (i.e., render them false), that's evidence that *e* is context sensitive.

Take the first person pronoun 'I'. Sentences containing 'I' cannot be disquotationally indirectly reported (except by self-reporters). If Bill tries to report Frank's utterance of 'I am Italian' with 'Frank said that I am tired', he'll fail. Utterances of sentence containing 'today' cannot be disquotationally reported (except by same-day reporters). If Bill tries today to report Frank's utterance yesterday of 'John will leave today' with 'Frank said that John will leave today', he'll fail. Utterances containing the demonstrative 'that' cannot be disquotationally reported (except by co-demonstrating reporters). If Bill tries to report Frank's utterance of 'That's lovely' – with Frank pointing at a table with his own utterance of 'Frank said that that's lovely', while pointing at a car, he'll fail. And so on.

In the reverse direction, suppose you suspect expression *e* is context *insensitive*. Take an utterance *u* of a sentence *S* containing *e* in context *C*. And let *C'* be a context relevantly different from *C* (e.g., different speakers, different times, different demonstrated objects, different speaker intentions, different salience conditions, and so on). If there's a true disquotational indirect report of *u* in *C'*, that's evidence *S*, and so *e*, are context *insensitive*.⁷

For example, consider three contexts *C1-C3* varying however you like. In each, you utter (1).

1. 'gone' is dissyllabic.

We'll now engage in actual speech acts: The indented items below represent actual utterances made by us in our study. These are acts where we are describing your three utterances of (1),

-In *C1*, you said that 'gone' is dissyllabic.

-In *C2*, you said that 'gone' is dissyllabic.

-In *C3*, you said that 'gone' is dissyllabic.

Our claim is that we can so report utterances of (1) in *C1 - C3* without remembering or even ever knowing any particulars about the original context of utterance. In particular, we need not have extensive knowledge of:

⁷ For further discussion of this test, see Cappelen/Lepore (2006; 2007).

- Speaker intentions
- Audience Intentions
- The nature of the conversation the speaker and audience were engaged in
- The assumptions shared among participants in the original context
- What was contextually salient in the original context of utterance
- The perceptual inputs of those participants in the original context.

If we're right about this, it provides very strong evidence that quotation is *not* (semantically) context sensitive.

We turn to our second test for semantic context sensitivity.

Collectivity Test: If a noun *e* is (semantically) context sensitive (i.e. if it changes its *semantic* value from one context of utterance to another), then on the basis of merely knowing there are two contexts of utterance, one in which 'e is F' is true and the other in which 'e is G' is true, we *cannot* automatically infer there is a third context in which 'e is F and G' is true.

As an illustration, suppose there are two contexts in which 'Yesterday John left' and 'Yesterday Bill left' are true (though we know little about these contexts; pertinently, we don't know the days of the contexts). We surely couldn't infer there is a third context in which 'Yesterday John and Bill left' is true. 'yesterday' does not admit of blind collection. The same holds for all of the other familiar context sensitive expressions. Test for yourself.

Returning to quotation expressions, recall García-Carpintero's examples (1) and (3).

1. 'gone' is dissyllabic.
3. 'gone' sounds nice.

According to him, the quotation expression in (1) refers to an expression and in (3) the same quotation expression refers to a different type somehow related to the token.

Suppose he's right. Here's evidence against a semantic construal of his data. Suppose you know no more than that there is a context C1 in which (1) is true and a context C2 in which (3) is true. Is it your intuition that there is a third context in which (12) is true?

12. 'gone' sounds nice and is dissyllabic.

To the extent that you share our intuition that this context exists, regardless of how ignorant, misinformed or disinterested you are in C1 and C2, then the more likely it is you will reject his claim that quotation is semantically context sensitive.

The bottom line is this: unprejudiced applications of the Indirect Disquotation Reporting and Collectively Tests (and others – see Cappelen/Lepore (2004, Ch. 7)), provide evidence against the claim that quotation is semantically context sensitive. We don't deny that if you come to the data with a philosophical axe to grind, you might very well be able to screen off the relevant intuitions; our point is that from a naïve perspective the reactions described above seem unimpeachable *and* nothing like this is possible for any of the traditionally accepted context sensitive expressions.⁸

The conclusions of Part I then are as follows: With respect to the semantic thesis, the legitimacy of the Strong Disquotational Schema and the (somewhat mysterious) proximity relation are inexplicable if QCS is construed semantically. Furthermore, quotational expressions regularly fail standard tests for semantic context sensitivity. The view that we can quote tokens also seems to have counter-intuitive consequences. On the assumption that this is all correct, we need an alternative explanation for the remaining variability data. We turn to that task directly.

Part II: Signs, Expressions and the Individuation of Quotables

The arguments in Part I establish the data in support of QCS do not establish that quotations are semantically context sensitive. The task in this part is to theorize about how best to accommodate data in support of QCS. We proceed in three steps: first, we outline three more or less familiar pragmatic strategies for explaining data about

⁸ It is surprising to us, to say the least, that we hadn't noticed problems with a semantic QCS sooner. Two research projects have occupied most of our philosophical attention during the past decade: the semantics of quotation (Cappelen/Lepore 1997a, 1998, 2005) and the limits of context sensitivity (1997b, 1999, 2000, 2004). What we failed to notice is that our commitments with the latter research program prohibit our commitments with the former. We embraced a Demonstrative Theory of quotation (Davidson 1979; Cappelen/Lepore 1997a), according to which quotations are treated as contributing a definite description with a *demonstrative* element, so the quotation marks in (1) abbreviate the demonstrative description 'The expression of which this is a token', as in (1*):

1. 'gone' is dissyllabic.

1*. Gone. The expression of which that is a token is dissyllabic.

As the Demonstrative Theory explicates quotations demonstratively, it is committed to a semantic QCS. García-Carpintero explicitly exploits this component to explain the variability data. What a demonstrative refers to varies, thus providing the flexibility to explain how the referent of "gone" can shift from one context to another (García-Carpintero 1994, p. 261). If you agree, however, with us that the Indirect Disquotation Report and Collectivity Tests (and others) limit claims about (semantic) context sensitivity and discredit a semantic QCS, then you must reject the Demonstrative Theory.

communicated content that is not part of semantic content and suggest how these strategies can account for some of the variability data. In §2.1, we reexamine some of Garcia-Carpintero's data that resist pragmatic explanation. In §§2.2-2.3, we develop an alternative syntactic framework for explaining this recalcitrant data.

§2 *Explaining the Variability Data*

Here are three familiar strategies for explaining intuitions about a sentence S varying in truth-value between contexts of utterance even though it is presumed that S isn't semantically context sensitive:

- **Strategy 1: Conversational Implicature:** In some cases cited in defense of a semantic QCS (see above), we could say what the speaker *said* is false, but what she *implicated* is true. For example, for Garcia-Carpintero's example (5) (that we already dismissed as a illegitimate case of semantic context sensitivity):

5. At least one of these words is heavier than 'gone'.

(which you are supposed to imagine written in big wooden letters), we might opt to read such that each of its utterances semantically expresses a false content (i.e. assuming the semantic value of 'gone' is abstract, and so, lacks weight) but, just the same a speaker can expect her audience, by virtue of recognizing its obvious absurdity, to conversationally infer that a contextually salient token has a weight.

- **Strategy 2: Speech Act Pluralism:** Alternatively, if you agree with us (and others)⁹ that an utterance of a sentence S (literally) asserts a plurality of propositions, only one of which is semantically expressed,¹⁰ then you can say about the cases in question that they literally assert truths. Since these truths need not be the semantic content of S, we need not infer a semantic QCS. On this view, it is possible for an utterance of a sentence S containing a quotation expression Q to (literally) make an assertion about a token, even though that is not the semantic content of S and Q.
- **Strategy 3: Pragmatic Ellipsis:** In some cases, an appeal to ellipsis might explain the data. An utterance of (5) might be elliptical for (5')

5'. At least one of these words is heavier than this token of 'gone'.

⁹ See Cappelen/Lepore 2004, Chapter 13.

¹⁰ Note that on some versions of speech act pluralism, you don't need to assert the semantic content. See Cappelen/Lepore 2004, Chapter 13.

How attractive this option seems to you will depend on what your favorite theory of ellipsis says.¹¹

Our goal in what follows is not to choose among these strategies. As we see the situation, the arguments of Part I rule out a semantic QCS. The question that remains is which non-semantic strategies are available to explain intuitions about the legitimate variability data. Strategies 1-3 are familiar and can be applied in rather obvious ways to explain some of these data. The goal of the remainder of this paper is to develop a fourth strategy – neither semantic nor pragmatic, but *syntactic*. Why bother given that Strategies 1-3 are available? Because some of the variability data we have presented in defense of QCS resist a pragmatic (and obviously semantic) interpretation; the persuasiveness of this residual data is attributable to a failure to individuate correctly quotation expression types, i.e., in order to see the data in question as supporting QCS, you need to mistake two utterances as being of the same sentence-type. Since we deem the utterances to disagree in truth-value, we mistakenly infer this is because a quotation expression/sentence is context sensitive. We'll argue that the correct conclusion (in some cases at least) is that the utterances in question are not of the same sentence type. We return first to the recalcitrant data that provoke this radical departure.

§2.1 Recalcitrant Data

A problem conveniently skirted in Part I is that some of the variability cases canvassed earlier permit *neither* easy indirect disquotational reporting *nor* collection of prior uses. We focus on such cases in what follows because they clarify the significance of quotation-expression individuation.

Reconsider Garcia-Carpintero's example (6) from above:

6. 'Madrid' = '*Madrid*'.

According to Garcia-Carpintero, there are contexts in which these quotation expressions have the same content, but others in which they do not. Imagine (6) tokened in a context C where the speaker is typing a rather obvious identity, but mid-sentence his italicization key jams forcing italics. In this context, (6) seems true. Next imagine (6) tokened in a context C' (C ≠ C') where the speaker is intending to bring to his audience's attention to differences between the items flanking the identity symbol. In this context, (6) seems

¹¹ Not in the syntactic sense, but in the sense of Sellars (1954, pp. 198ff) and Neale (2000, p. 187).

false. The problem is that these data seem to flout our tests for semantic context insensitivity. They permit neither of collection nor indirect disquotational reporting.

Try collecting the imagined true token of (6) in C and the imagined true token of its negation (13) in C' into a true token of (14), say, in C'':

13. 'Madrid' ≠ *'Madrid.'*

14. 'Madrid' = *'Madrid'* and 'Madrid' ≠ *'Madrid'*.

You can't do it without seemingly invoking incoherence.

As for indirect disquotational reports, try to report a written token of (13) (uttered in C') in spoken language (i.e. let the indirect report of the written utterance of (13) in C' be spoken) in a context C''. You run into obvious difficulties. A spoken medium doesn't permit us (semantically) to express by mere disquotation of (13) what's going on in these sorts of cases. Or, even suppose your keyboard lacks italicization. It would seem that you couldn't even indirectly disquotational report a token of (13) graphemically.

How are we to explain (away) these puzzling data? A semantic story is off the table for reasons provided in Part I. Invoking pragmatic strategies 1-3 is not going to help much in these contexts either. Suppose we try to argue, for example, that one of (6) or (13) is false. If so, which? And what would be the reason for it and not its counterpart being false? Or suppose we try to argue that, contrary to intuition, a spoken indirect disquotational report of (13) is as a matter of fact true. The extreme implausibility of this view makes such a position exceedingly unattractive. Similarly, suppose we try to argue (6) is false. If so, are (15) (where the two quotation expressions are in different fonts) and infinitely many like identities also false?

15. 'Madrid' = 'Madrid'.

Is (16) (where its two quotation expressions are written in the same font but in different sizes) also false?

16. 'Madrid' = 'Madrid'.

Obviously, what's needed, and what has been transparently lacking in the discussion thus far, is an explanation that addresses issues about how quotation expressions and quotable items are individuated. Are the two quotation expressions in (6) identical or distinct?

Quotations themselves are linguistic expressions; otherwise, they could not be legitimate constituents of sentences. And indeed, every quotation sentence sampled in

this paper has a quotation as its grammatical subject or object. Anyone who endorses a semantic or pragmatic interpretation of QCS based on variability intuitions, therefore, assumes these intuitions are about distinct utterances, for example, of (1)-(5) are utterances of the *same* quotation-expression. This is the assumption we will challenge (at least for some cases) in what follows. To do so, we have to answer two questions:

D1: How could the utterances in question (e.g. of (1)-(6) in C and C') be of different quotation-expressions?

D2: Why have theorists and informants thought otherwise, i.e. why do theorists and informants tend to believe that these utterances are of the same quotation expression?

We begin with D1.

§2.2.1 *1st Step towards answering D1: Signs*

How can we discuss the puzzling variability data without presuming expression type-identity across distinct contexts of utterance? To answer this question, we digress into a brief discussion about *signs*. Its significance to expression-type identity, in particular, to quotation-expression-type identity, will soon become apparent.

Different media can obviously be used to articulate the same message. We can write English; we can speak it; we can write it with the Roman alphabet (and in indefinitely many fonts, sizes), in italics, in bold, in Braille, in Morse code; we can sign it or wave it in semaphore; we can speak it with different pronunciations. These constitute but a tiny fragment of indefinitely many distinct ways in which we can articulate the very same language, the very same expressions. Call each of these a *sign system for articulating English*.

No language is essentially tied to any particular sign system or even to any set of sign systems. Though it might be that for most natural languages spoken sign systems take on more practical significance than other sorts of sign systems, if as a result of a global catastrophe, the human race goes mute, we could still read this sentence, and you would be reading an English sentence. You could still write or record a reply to it, and that reply could be in English. Our point is that no single sign system or set of sign systems is essential for the continued existence of English or for any other natural language for that matter.

As a result, speakers who do not share a sign system can still share a language. Consider A, B and C. A speaks, but is illiterate (i.e., does not write or read). B is deaf/mute and only reads and writes. C is blind and deaf and only communicates in Braille. A, B and C can all be English speakers, even though they are incapable of communicating for lack of a common sign system. And in the other direction, obviously people who share a sign system may not share a language. Our friends in Italy write to us using the same sign system that we do, but they use it to articulate Italian words and we use it to articulate English words. It's convenient for us to use the same sign system; indeed, our common computer keyboards and software insist upon it¹².

§2.2.2 2nd Step towards answering D1: What are the quotable items?

As it goes with any expression type, a quotation expression when articulated must be articulated with a sign. That part is uncontroversial. A sign, like any other kind of entity, can be named; it can be described or demonstrated; and once picked out, all sorts of attributes can be ascribed to it. We might discover of a particular sign that it is used in articulating different expressions in one or many languages – expressions which differ in meaning, and so on and so on. These we presume are also all uncontroversial points.

In addition to being namable, describable and demonstrable, a sign is also quotable, as in (17).

17. 'red' is a sign which articulates one word in English and a different word in Norwegian.

According to (17), one and the same sign is used to articulate different words (in different languages). (17) requires for its truth that its quotation expression (it's grammatical subject) has as its semantic value, that is, has as whatever item it quotes, a particular sign. The semantic value of the subject of (17) can't be two different words; rather it's a single sign.

¹²A theory of signs needs to say something about how they should be individuated. Must two items in different fonts be different signs? What about differences in font size? Do all handwritings constitute different sign systems? What about different pronunciation patterns? We will not attempt to answer these extremely difficult and interesting questions here. Quite frankly, there are a number of issues about sign systems we are not quite sure about at this point that we hope to explore in future work. What we are quite sure about is that these issues will not affect the *structure* of the view we defend in this paper. We should also note that there's an important connection between how one thinks of the individuation of signs and questions about how many categories of quotables there are; see nt. 16 for further remarks about this.

In addition to signs being quotable items, obviously, expressions are as well, as in (18):

18. 'red' is an English word.

Notice that we can say of this word quoted in (18) that in our spoken sign system it's pronounced the same way as the word 'read' is. So, we have distinct graphic signs for articulating two different words, but for some historical reason we have the same spoken sign for articulating them both. Again, we presume all of this is obvious.¹³

In sum: differences between (17) and (18) indicate that quotation expressions can quote distinct kinds of items. What's quoted in (18) is from a certain language (English), it has a certain meaning. Signs, on the contrary, are not language specific; they lack meaning (except possibly in some derivative sense). We don't infer that the word 'red' is ambiguous because two words with distinct meanings are pronounced in the same way (i.e. with the same spoken sign).

That we are committed to there being different sorts of quotable items – signs *and* expressions¹⁴ – doesn't, of course, distinguish us from anyone who endorses either a semantic or pragmatic construal of QCS: semantic and pragmatic variability *only* makes sense on the assumption that a quotation expression can be used to quote distinct sorts of quotable items.

§2.2.3 3rd Step towards answering **DI**: Containment

Since both signs and expressions are quotable items, if quotation expressions are unambiguous or context insensitive, it follows that if a quotation expression quotes an expression, it cannot be quoting a sign and if a quotation expression quotes a sign, it cannot be quoting an expression. On this view, (19) and (20) can both articulate true English sentences.

19. 'red' is a sign and not an expression.

20. 'red' is an expression and not a sign.

In the first, the grammatical subject quotes a sign, and in the second the grammatical subject quotes an expression. How, given the arguments of Part I, are we to explain this possibility? If you re-examine the strong disquotational schema SDS, you'll see – as we

¹³ The distinction between signs and expressions in the quotation literature goes back at least to Linsky (1950) and perhaps also to Reichenbach (1947).

¹⁴ For some important qualification, see nt. 16 below.

flagged at the time – that according to it quotation expressions *contain* as constituents their semantic values – namely, the quotable item is itself a constituent of the quotation expression that quotes it. Although we initially spoke of this notion of containment in somewhat metaphorical terms, we intend it literally. The sense of ‘containment’ we have in mind here is intended to be no less innocent than its sense when we say of a sentence that it contains any number of expressions.

If quotation expressions contain whatever they quote as constituents, it follows that if a quotation expression quotes a sign, then that sign is a constituent of that quotation expression; and if, on the other hand, a quotation expression quotes an expression, then that expression is a constituent of that quotation expression. In short, two quotation-expressions are distinct just in case they disagree in semantic value (or much the same, just in case they don’t share identical constituents).

The relevance of these claims for context sensitivity should be clear. The tests we advocated for discerning semantic context sensitivity – the Indirect Disquotational Reporting and the Collectivity Tests – both presume linguistic stability – by which we mean both presume that what’s indirectly disquotationally reported or collected are the same expressions as those originally tokened. Likewise, pragmatic strategies 1-3 presume linguistic stability; each assumes that the same quotation sentence might convey one proposition in one context and a different one in another context. However, on the current proposal for individuating quotation-expressions, a failure to collect, for example, (6) and (13) into (14), is a result of a *syntactic*, and *not* a semantic, equivocation. For the two imagined true utterances in question, the subject of (13) is *not* the same expression as the subject of (6):

6. ‘Madrid’ = ‘*Madrid*’.

13. ‘Madrid’ ≠ ‘*Madrid*’.

14. ‘Madrid’ = ‘*Madrid*’ and ‘Madrid’ ≠ ‘*Madrid*’.

Since their grammatical subjects are not the same neither are their subject matters, and so trying to collect these two subjects into a single sentence is as confused as trying to collect expressions from different languages into a single sentence in a single language. No one would conclude an English word is ambiguous or semantically context sensitive because a sign that articulates it in English also articulates other words from other

languages with different meanings. Likewise, were you to try to indirectly disquotationally report a true token corresponding to the non-identity in (13) you might fail were you to report it out loud – not because of semantic context sensitivity but because spoken signs are distinct from written ones, and so, if signs are being quoted then obviously the one medium cannot be utilized in reporting the other – if the medium itself provides the subject matter of the quotation sentence.¹⁵ However, were you to try to indirectly disquotationally report a true token corresponding to the identity in (6) you might succeed since the quoted items might be the same word, and not two distinct signs.

Since both semantic and pragmatic tests for context sensitivity presume that the same expression is being indirectly disquoted or collected in their applications, and since, on the current proposal, that presumption is false for the recalcitrant variability data, the data is explained (away).

§2.2.4 *Summary of answers to D1*

Which expression is employed and which sign is used in articulating that expression are contingent upon speaker intentions and, perhaps, contexts of usage. But the way in which these intentions (or contextual salience) determine which sign system and which language a speaker is employing on a given occasion is *not* the same as how speaker intentions and contextual salience determine the semantic values of context sensitive expressions, or even conversational implicatures (or any other speech act exceeding semantic content). Nor is it pre-semantic, in Kaplan's sense. Kaplan says:

Given an utterance, semantics cannot tell us what expression was uttered or what language it was uttered in. This is a pre-semantic task. (Kaplan, 1989, p. 559)

Similarly, if a word is ambiguous, semantics won't say which meaning is intended in a particular utterance. The effect a pre-semantic process has on semantic content is 'indirect'. Even though language selection influences the semantic content of the utterance, it does so indirectly through the determination of the input to the semantic machinery that churns out semantic content. And the choice we must make when

¹⁵ Once signs are recognized as quotable items it should be obvious that there will be expressions, namely, some quotation expressions, that can be articulated in one medium but not in others – a case of the tail wagging the dog.

confronted with ambiguity has the same pre-semantic, hence ‘indirect’, effect on semantic content.

In this regard, a speaker’s choice of a sign system is *not* even pre-semantic. *After* disambiguation and language selection have occurred, a choice still remains of which sign system to exploit on an occasion of articulation, but this choice has no impact – direct or indirect – on semantic content. For suppose P is a proposition that A wants to express in a language L with a sentence S on a given occasion. He can articulate S only after choosing a sign T. However, T must articulate S in order to express P or else A has made a mistake. Of course, an element of context sensitivity remains (though not semantic): For which sign system is used on a given occasion is itself determined in context. But this is true no matter how you twist and turn the data, so it should come as no surprise.

These claims are a good first (and independently motivated) effort at explaining (away) the recalcitrant variability data. They obviously require further development – more than we can hope for here, but we can go further; for we can explain *why* theorists (including former time-slices of ourselves) and informants have missed this explanation of the data, i.e. we can answer D2. In the next section we propose three such explanations.

§2.3 *Answers to D2*

There are many reasons why theorists and informants fail to distinguish sentences with distinct quotation expressions (i.e. why they’re blinded to the nature of quotable items). There are various hurdles an interpreter must get over in order to correctly identify a particular quotation expression. We’ll survey some in what follows (though we don’t intend our list to be exhaustive). Failure to get past any can easily lead to a misidentification of a quotation-expression, and so, as to what it quotes.

§2.3.1 *From Tokens to Signs*

Each interpreter begins with a physical token. His first task is to figure out which sign this physical token instantiates. This problem does not admit of an easy solution.

First of all, two tokens might be physically indistinguishable and yet still constitute distinct signs. A sign is the sign it is not merely because of its intrinsic physical properties but also by virtue of its relationship to other signs in a sign system. In effect,

sign systems are sets or bundles of contrastive properties where each member of a sign system differs from every other member at least with respect to one of these features. This of course is compatible with two tokened signs from different sign systems being physically indistinguishable.

Secondly, in the other direction, physically distinguishable items might be the same sign. Suppose two speakers send emails to each other using figures of different sizes. These physical differences might not matter for the purposes of exchange and so may not suffice for distinguishing signs and sign systems. In this regard, with respect to some sign system, an 'a' and an 'a' may be the same sign – though physically distinct.

Returning to Garcia-Carpintero's own example (6),

6. 'Madrid' = '*Madrid*'

whereas he thought the variability data surrounding (6) established that it is semantically context sensitive, we want to conclude that one utterance might correspond to a trivially true identity between the same quotation expression and the other to an obviously false identity between distinct quotation expressions. As already noted, in order to determine which it is, a speaker must first determine whether physically distinct symbols flanking a tokened identity sign are the same signs or not.

This shouldn't be too surprising; after all, no two of us write in exactly the same style – our handwritings are distinguishable. People make a career out of being able to pin a person to a scrawl. The same goes for our pronunciations. Actors put a great deal of effort into mimicking regionally displaced speakers of the same language. Yet no one wants to infer on these bases alone that the signs employed in any two distinguishable handwritings are distinct. The output of pencils, pens, chalk, keyboards, and other sign producing devices might be physically distinct without their signs being so.

A second problem we encounter in deciphering another's message arises inasmuch as we are assuming quotation expressions can quote (at least) signs and expressions.

§2.3.2 *From Signs to Expressions*

The problem touched on in the last subsection concerns the individuation of signs – a topic too big for this occasion but one whose subtleties are easy enough to grasp, especially in the context of explaining why we might misidentify a quotation expression

when confronted with its token. Notice that since expressions are once removed from the signs that articulate them, when a quoted item is an expression and not a sign, the problem of quotation expression identification is only compounded.

For even after correctly identifying the signs in play with regards to any articulated quotation sentence, we still must ascertain whether its quoted item is the sign used in articulating the quotation expression or the expression that that sign articulates (if any).

Once, again, returning to Garcia-Carpintero's example (6),

6. 'Madrid' = '*Madrid*'

as he describes the data surrounding it, we are asked to imagine a true utterance of (6). Our preferred way of putting the request is that, having identified the sign used in articulating two quotation expressions flanking an identity, we next need to determine which expressions they are articulating. Either identical or non-identical signs can be used to articulate distinct quotation expressions – one sign might articulate a quotation expression that contains a sign as its constituent and another might articulate a quotation expression that contains an expression, say, the name for the capital of Spain, as its constituent. Often, when the same sign is used to articulate expressions that differ in meaning we posit ambiguity – i.e., the expressions are the same; just the meanings differ. But with quotational expressions this maneuver is unavailable. This is because if the quotational expressions that the identical signs articulate have distinct semantic values, it follows, on our syntactic account, that the expressions themselves are distinct (because their constituents – their semantic values – are distinct). So even having identified the signs employed on a given occasion, we still have our work cut out for us – we still need to determine which expressions are being articulated. This extra work provides us with yet another opportunity for failure.

§2.3.3 *Quotation across Sign Systems*

Once you recognize the differences between signs and the physical types their tokens instantiate, and also between expressions and signs used in articulating them, some of the confusion surrounding the recalcitrant data becomes transparent. We can articulate, as we have noted several times, the same expression in indefinitely many sign systems. We can therefore quote the same expression in indefinitely many sign systems.

But we *can't* quote the same sign in indefinitely many sign systems (cf. nt. 12). Indeed, it's unclear whether a sign from one sign system can be quoted in any other sign system. According to SDS, whatever is quoted must be contained in the quotation expression that quotes it. So if a sign is quoted by a quotation-expression, then that sign must be contained in that quotation-expression. We also assumed earlier that what individuates a sign within a sign system are certain contrastive features. If these differ between sign system T1 and T2, no sign from T1 can be quoted in T2. To that extent, sign quotation is sign system bound.¹⁶

§2.3.4 *Summary of answer to D2*

Differences can result when distinct signs are exploited and/or invoked. However we conceive of sign individuation, it is not the same as ambiguity or semantic context sensitivity, that is, differences in meaning are not driving this discussion. In the different cases described, the same quotation expression is not employed throughout. Rather, in the recalcitrant cases, distinct quotation expressions discuss different quotable items. Differences in the signs used are (partly) responsible. Differences in intuited truth-value do not reflect semantic context sensitivity, then, since no expression has its semantic value determined by its context of utterance. Differences in sentential truth value are due to differences in sentences (and thus expressions) articulated.

Conclusion

Where have we gotten to? We concede that different information might be conveyed by the same expressions on distinct occasions of use, and so, we understand why someone might mistakenly conclude that a given expression is semantically context sensitive. However, we argued that quotation is not semantically context sensitive. For one, if it were, we would be forced to relinquish something sacred about quotation –

¹⁶ Why assume that nothing more basic than signs and expressions can be quoted? If there are true sentences of the form '...' is not a sign in any sign system or an expression in any language' (where some quotable item replaces the three dots), then there must be quotable items that exist outside sign systems and languages. It may be that non-philosophical speakers use of quotation support the hypothesis that there are such items (we consider that question up to further investigation). What is crucial to note is that if the class of quotable items turns out to contain kinds of entities other than signs and expression, the structure of the view defended in this paper is preserved. According to the containment principle invoked in §2.2.3, any quotable item must enter into the individuation conditions for quotation expressions in the same way that signs and expressions do. That's the principle invoked by the syntactic explanation of the recalcitrant data. So, we wouldn't be too bothered were it to turn out that the metaphysics of quotable items incorporates more than signs and expressions.

namely, its disquotational nature. To ensure its disquotational nature, we proposed a rather obvious and simple semantic rule for quotation – namely, SDS. A consequence of accepting SDS is the surprising result that quotation expressions contain their semantic values as constituents. Once that view is embraced, we provided a fourth, syntactic, strategy for explaining the appearance of context sensitivity for quotation that went beyond the three pragmatic strategies we exploited in explaining some of the variability data. The key to our solution for the recalcitrant data is to see that quotation expressions can quote not only expressions but also signs. The distinction is well evidenced independently of any considerations about the theory of quotation. We ended with a series of explanations for why theorists have failed to recognize that physically indistinguishable tokens might yield distinct quotation expressions.

We wound up, we believe, with a rather elegant story about the nature of quotation, but work remains to be done. As we see it, the proposed solution raises a number of important questions that require further investigation, some of which are:

- What are sign systems and how are they individuated?
- Do signs and expressions exhaust the set of quotable items?
- If expressions are distinct from sign, what kind of objects are expressions?

As we point out above (in note 16), we predict that answers to these questions will preserve the structure of the solution presented here, as long as the containment principle invoked in §2.2.3 is preserved. In that sense, the theory of quotation presented here is amenable to a rather wide range of implementations.¹⁷

¹⁷ Our views on quotation seem to be in perpetual evolution – not only since our initial publication in 1997 but even since the conception of this very paper. The view articulated in our first presentation at the Society for Exact Philosophy meeting in Toronto and the University of Santiago de Compostela last spring (2005) is quite different from the view articulated in our presentations of the view this past fall (2005) at the Universities of Kansas and Oslo. And this latter view is yet still distinct from the one present in this draft of the paper and most recently presented at the University of Florida and at the University College Dublin.

We'd like to thank all these audiences for participating in the evolution of our view, especially John Hawthorne – who has provoked us at every step along the way. It's a bit frustrating to change our minds about a topic we have thought about as long as we have thought about quotation, but we must always strive to do better.

Bibliography

- Atlas, J. 1989. *Philosophy Without Ambiguity*, Oxford University Press, Oxford..
- Cappelen, H. and E. Lepore, 1997a. 'Varieties of Quotation'. *Mind* 106, pp. 429-50.
- -----. 1997b. 'On an Alleged Connection between Semantic Theory and Indirect Quotation', *Mind and Language*, vol. 12, pp. 278-296.
- -----. 1997b. 'Varieties of Quotation'. *Mind* 106, pp. 429-50.
- -----. 1998. 'Using, mentioning, and quoting: Reply to Tsohatzidis'. *Mind* 107, pp. 665-666.
- -----. 2004. *Insensitive Semantics*. Oxford: Basil Blackwell Publishers.
- -----. 2005. 'Varieties of Quotation Revisited', *Belgian Journal of Linguistics*. 17, pp. 61-76.
- -----. 2006. 'Reply to Hawthorne' *Philosophy and Phenomenological Research* forthcoming.
- -----. 2007 'On the Semantic Significance of Disquotational Reports' in *Context-Sensitivity and Semantic Minimalism*. (eds.) G. Preyer and G. Peter, Oxford University Press: Oxford, forthcoming.
- *Language Turned on Itself*, in progress.
- Christensen, Niels. 1967. 'The alleged distinction between use and mention'. *Philosophical Review* 76, pp. 358-67.
- Clark, Herbert & Richard Gerrig. 1990. 'Quotations as demonstrations'. *Language* 66(4), pp. 764-805.
- Davidson, D. 1979. 'Quotation', in *Inquiries Into Truth and Interpretation*, Oxford, Oxford University Press, 1984, pp.79-92. Originally published in *Theory and Decision*, 1979 11, pp. 27-40.
- Garcia-Carpintero, Manuel. 1994. 'Ostensive signs: against the identity theory of quotation'. *Journal of Philosophy* 91, pp. 253-64.
- -----. 2004. 'The Deferred Ostension Theory of Quotation'. *Noûs* vol. 38, no. 4, pp. 674-692.
- -----. 2005. 'Double-Duty Quotation: The Deferred Ostension Account', *Belgian Journal of Linguistics*, pp. 89-108,
- Goddard, L. and R. Routley. 1966. 'Use, mention, and quotation'. *Australasian Journal of Philosophy* 44, pp. 1-49.
- Goldstein, Laurence. 1984. 'Quotation of types and types of quotation'. *Analysis* 44, pp. 1-6.
- Gomez-Torrente, Mario. 2001. 'Quotation Revisited'. *Philosophical Studies* 102, pp.123-53.
- -----. 2005. 'Remarks on Impure Quotation', *Belgian Journal of Linguistics*, pp. 129-152.
- Jorgensen, Julia, Miller, George and Sperber, Dan. 1984. 'Test of the Mention Theory of Irony'. *Journal of Experimental Psychology: General*, 113, pp. 112-20.
- Kaplan, D. 1989. 'Demonstratives'. In *Themes from Kaplan*, eds. Almog, J., J. Perry, and H. Wettstein, Oxford University Press: Oxford, pp. 481-564.
- Linsky, Leonard. 1950. 'On Using Inverted Commas'. *Methodos* 2, pp. 232-236.

- Ludwig, K. and G. Ray. 1998. 'Semantics for Opaque Contexts.' *Philosophical Perspectives* 12, pp. 141-166.
- Mates, B. 1972. *Elementary Logic*, 2nd edition, Oxford: Oxford University Press.
- Neale, S. 2000. 'On Being Explicit: Comments on Stanley and Szabo, and on Bach'. *Mind and Language*, 15, pp. 284-94.
- Quine, W.V.O. 1940. *Mathematical Logic*, Boston, Mass., Harvard University Press.
- Reichenbach, H., 1947. *Elements of Symbolic Logic*, Free Press: New York.
- Reimer, Marga 1996. 'Quotation Marks: Demonstratives or Demonstrations?' *Analysis*, 56, pp. 131-42.
- Richard, Mark 1986. 'Quotation, Grammar, and Opacity'. *Linguistics and Philosophy*, 9, pp. 383-403.
- Saka, Paul 1998. 'Quotation and the Use-Mention Distinction'. *Mind*, 107, pp. 113-35.
- ----- 1999. 'Quotation: a reply to Cappelen & Lepore'. *Mind* 108, pp. 751-54.
- ----- 2005. 'Quotation constructions'. *Belgian Journal of Linguistics*.
- Salmon, N., 1986. *Frege's Puzzle*, Cambridge, Mass: MIT Press.
- Sellars, W. 1954. 'Presupposing'. *Philosophical Review*, 63, 197-215.
- Smullyan, R.M. 1957. 'Languages in which Self Reference is Possible', *Journal of Symbolic Logic*, 22, pp.55-67.
- Wallace, J. 1972. 'On the Frame of Reference', in *Semantics of Natural Language*, eds., D. Davidson and G. Harman, pp. 219-252.
- Washington, C. 1992. 'The Identity Theory of Quotation', *Journal of Philosophy*, 89, pp. 582-605.
- Wertheimer, R. 1999. 'Quotation apposition'. *Philosophical Quarterly* 49(197), pp. 514-19.