

*Meaningfulness and Contingent Analyticity**

ORI SIMCHEN

University of British Columbia

Actually sentences like ‘Socrates is called “Socrates”’ are very interesting and one can spend, strange as it may seem, hours talking about their analysis. I actually did, once, do that. I won’t do that, however, on this occasion. (See how high the seas of language can rise. And at the lowest points too.) (Kripke 1980, 73)

1. Introduction

As philosophers of language we are drawn by two aspects of the meaningfulness of linguistic expressions that tug us in seemingly opposite directions. The first aspect stems from the fact that the items we employ in linguistic exchange are ultimately physical items—utterances and inscriptions. Notoriously, it can seem puzzling that such items possess the contents they do, for it can seem puzzling that otherwise lifeless physical items can be animated with content at all. But whatever we are inclined to think about this thorny issue, the fact that a particular linguistic expression, regarded as a physical item, should possess the content that it does can only be a contingent fact about it. Surely nothing about the item itself necessitates that it should have whatever content it happens to have.

The second aspect of meaningfulness which seems to conflict with the first stems from what may be regarded as the linguistic closure of reflection. All higher reflection is conducted in language. So the fact that linguistic expressions should in general possess the contents they do is incontestable to the extent that contesting is a species of higher reflection. Being a species of such reflection, it is itself a linguistic activity. And in order for it to be achieved, the expressions used are required to have their contents. Moreover, if a specific doubt as to the content of a given expression involves

using the expression itself, the doubt appears to undermine itself. In order for the doubt to be achieved, the expression in question is required to possess the contested content. In this way it can seem a priori true that an expression should have the semantic content that it has. Claiming that it has it involves using the expression and claiming of it that it has the content it already must have within this very claim.

This paper aims to resolve some of the apparent conflict between these intuitions about meaningfulness: that it is contingent and that it is a priori. The conflict has implications for some of the central issues in the philosophy of language. To pick a prominent example that won't be discussed here, Tarskian T-sentences have been viewed in two distinct ways. On the one hand, and most famously, they have been regarded as serviceable within an empirical science of semantics as criteria for the correctness of the semantics. In order for the semantics to be correct, T-sentences expressing contingent matters of fact must be theorems of the theory. On the other hand, homophonic T-sentences such as "Snow is white" is true if and only if snow is white' are commonly thought to be true a priori. As such they are thought to be usable as purely formal criteria for the correctness of a semantics given in a metalanguage which extends the object-language.

Instances of disquotation of nouns such as "'Socrates" refers to Socrates' and "'cat" refers to cats' exhibit this ambiguity most clearly.¹ On the one hand they involve mentioning an expression that is being used in exactly the way it is claimed to be used. Thus, in a recent essay on Hilary Putnam's model-theoretic argument Bas van Fraassen claims of "'cat" refers to cats' that it is one among 'paradigmatic examples of pragmatic tautologies in [his] language' and goes on to say that it is 'undeniable by [him], exactly because [he] acknowledge[s] "'cat" to be a word in [his] language' (van Fraassen 1997, 42). However, there is clearly nothing about a mentioned expression that necessitates that it be employed as it is in fact employed. Accordingly, van Fraassen notes that '[t]he semantic content . . . of these . . . assertions are not necessary propositions, and most certainly not tautologies in the sense of semantics' (Ibid.). A similar observation is offered by Jonathan Harrison in a recent essay on Tarski: 'I cannot . . . state the fulfillment of the possibility that I am mooted when I consider that Henrietta's name might not be "Henrietta" in the words *Henrietta's name is not "Henrietta"*' (Harrison 1998, 9). 'But', Harrison hastens to add, 'this does not mean that, if such possibilities were realized, they would not be expressible at all'.

My main concern in what follows is to offer semantic and metasemantic underpinnings for this dual-aspect nature of disquotation.² The result will be an account according to which instances of disquotation, while expressing contingently true contents, are nevertheless analytic. My proposal will consist in bringing pivotal aspects of David Kaplan's framework of demonstratives to bear on a demonstrative approach to quotation.

A feature of Kaplan's framework that has received a good deal of attention in recent years is its accommodation of contingent analytic truths—sentence-types tokens of which express contents that are guaranteed by the characters involved to be true contents but not necessarily true.³ Familiar examples here include 'I am here now', 'I exist', 'You are my addressee', and 'It is now today'. The contingent analyticity in such cases crucially depends on the presence of indexicals—context-sensitive expressions for which the rules of the language are thought to fully specify the content in context. However, as I will show below, it turns out that sentence-types involving demonstratives, for which the rules of the language do not fully specify the content in context, may manifest the same contingent analyticity as those involving only indexicals, provided that we consider them along with types of determination of referent for each of their demonstratives. These are the demonstrative contingent analytic truths (henceforth, D-CATs). Thus, it so happens that contingent analyticity comes in two varieties: a purely indexical variety and a demonstrative variety. Once we see quotation as involving demonstrative reference to tokens, as I take it we should, the existence of D-CATs becomes the key to understanding how tokens of instances of disquotation, while not expressing necessary truths, are nevertheless guaranteed to come out true in every context of use.

Towards the end of the paper I turn to draw broader epistemological implications from my construal of disquotation by bringing it to bear on Putnam's brains in a vat (henceforth, BIV) scenario. In a cleaned-up version of Putnam's argument, the anti-skeptical response to the BIV skeptic depends on the observation that instances of disquotation are analytically true. Viewed in terms of my proposal, the skeptic advancing the BIV hypothesis is seizing on the contingency of the fact that an expression-token should have the reference it has. And the anti-skeptical response consists in pointing out that the skeptic underplays the analyticity of instances of disquotation. My construal of disquotation affords a clearer view than has previously been available of the encounter between the skeptic and the anti-skeptic. It reveals the anti-skeptical argument as attempting to rebuke BIV skepticism by appealing to the Kaplanian characters associated with types, and the skeptical response as shifting the terms of the debate to a consideration of the contents of tokens. But before any of this can be achieved, I need to present the proposal itself. And doing so requires turning, before all else, to the semantics of quotation.

2. Quotation

The most salient feature of quotation is usually taken to be that it creates an opaque context within the overall sentence—a context where even substitution of synonyms for synonyms, let alone non-synonymous co-referring expressions, may alter the truth-value of the whole. This makes quotation a

recalcitrant phenomenon for compositional semantics. Accordingly, the question of how to construe the semantics of quotation has been a matter of some dispute. A demonstrative approach to the problem was originally devised by Davidson in reaction to two then-influential accounts, the name theory and the description theory, which have proven unsatisfactory.⁴ This approach has recently been extended and refined, and has gained some prominence. Davidson takes quotation to be a means for referring to orthographic or phonetic shapes by exhibiting samples of those shapes, that is, inscriptions or utterances. Thus, for example, the direct discourse attribution,

(1) Marx was the first to say, ‘Workers of the world, unite’

is analyzed as

(2) Marx was the first to employ the shape of which that is a token.
Workers of the world, unite,

where the demonstrative ‘that’ of the first item refers to the tokens comprising the second item. Quoting, on this view, consists in mentioning linguistic types understood as shapes, via demonstrative reference to tokens.

This account makes vivid why quotation should create referential opacity. What we are attributing to Marx in (2) is being the first to employ the shape of which the second item of the analysis is a sample. And clearly changing the sample by substituting, say, ‘get together’ for the original ‘unite’, would make all the difference in what is being attributed to Marx. In other words, the explanation Davidson offers for the opacity of the quoting context is just that whatever is enclosed within quotation marks in cases like the one above is not a working linguistic component. It is a mere demonstratum. What does the requisite referring are the quotation marks and not whatever is enclosed therein or the quotation taken as a whole.⁵

Notice, however, what would have to be the case in order for (1) to come out true in light of its analysis (2). At least the following two conditions would have to be met. Marx would have had to employ a token of an English sentence-type, so presumably he would have had to coin his phrase in English. And he would have had to do so in writing. Since there is every reason to suppose that Marx did not coin his phrase by writing the English words ‘Workers of the world, unite’, our direct discourse attribution would most likely turn out to be false in light of its analysis. Yet there are clearly contexts in which we would be inclined to regard an attribution such as (1) as true. It is often simply obvious that we are quoting without regard to the particular language originally employed, let alone to the physical medium in which the words were employed. Think, for example, of (1) delivered as part of a lecture on the history of ideas, or appearing in a text-book on the

subject. That Davidson's theory should render such uses of quotation false is a marked disadvantage. Everyday uses of quotation vary greatly in specificity when it comes to what is being attributed to speakers. Sometimes we quote by way of attributing words in a particular language, unlike (1) in the history-of-ideas context. Other times we quote by way of attributing sound-patterns or shapes, as we might when we fail to understand what was said by their employment.⁶ And other times still we quote by way of ascribing the speaking of words with a particular intonation, in which case the intonation would be reproduced along with the quoted material. As it stands, Davidson's demonstrative approach to quotation cannot accommodate this rich diversity.⁷

But a recent elaboration on the demonstrative theory, due to Cappelen and LePore, can.⁸ Elsewhere I have argued that this demonstrative approach to quotation fares better than all of its competitors on several grounds, the most significant of which is its unique ability to handle cases of mixed quotation.⁹ On this improved version of the demonstrative theory, quoting includes demonstratively referring to a token that bears a certain relation, sametokening, to the token originally employed. Sametokening is a relation among tokens that is meant to remain as free as possible from any a priori constraints. We may take it to be analogous to Davidson's samesaying relation in the following way. Just as samesaying is supposed to obtain between tokens that would yield the judgment of synonymy that *what* was said by their employment on the relevant occasions was the same, sametokening may be said to obtain between items that would yield the judgment that *how* these items are is the same in the relevant respects. Whether or not tokens sametoken one another is thus a highly context-sensitive affair. Radically distinct items, such as phonetic and orthographic tokens, may count for specific quoting purposes as standing in this relation to one another. And sametokening may occasionally cut across languages, as, for example, in the case considered above concerning Marx, who did not coin his famous phrase in English, in the history-of-ideas context. Thus, (1) may now be construed as attributing to Marx being the first to employ a token which bears the sametokening relation to whatever is flanked by the quotation marks:

- (3) Marx was the first to say something sametokening that. Workers of the world, unite.

And in the relevant context (3) can indeed come out true. Notice once again that whatever appears as the second item of the analysis does not partake in the semantics of the whole qua linguistic item but only qua demonstratum.

3. Indexicals and Demonstratives

Given a demonstrative approach to quotation, we need to delve deeper and ask how we are to account for the demonstrative component of the analysis.

And here is where direct reference enters the picture. As is well-known, in the 1960's and early 1970's the direct reference revolt was launched from various sectors of the philosophy of language against the Fregean idea that linguistic expressions invariably express their senses, which, in turn, determine their referents. While Donnellan, Kripke, Marcus, and Putnam were concerned to unsettle this picture as it pertains to all proper names and most common nouns, Kaplan and Perry were concerned to unsettle it with respect to context-sensitive expressions such as indexicals ('I', 'you', 'yesterday', etc.) and demonstratives ('that', 'this', etc.). The basic innovation in this latter area, which has since become an orthodoxy, is a two-tiered picture of linguistic significance which may be sketched as follows.¹⁰

As against the Fregean distinction between sense and reference, it is now claimed that expressions are associated with their characters, which are rules specifying contents in contexts of use. Characters are, among other things, what competent speakers can be said to know in having command of the language. Context-insensitive expressions are associated with constant characters assigning the same content in every context, where the content in non-directly referential cases is a property, represented in the usual way as a function from circumstances of evaluation to extensions (an intension). But a context-sensitive expression such as the first-person pronoun 'I' is dramatically different. Its character is not constant, and the content assigned to it in context is not a property but an individual. Thus it is maintained that indexicals and demonstratives are directly referential, that they refer without the mediation of a property. And the difference between indexicals and demonstratives resides in the fact that in the former case the rules of the language fully determine the referent in context, whereas in the latter case an accompanying determination of the referent is required. We will come back to this difference in much more detail below.

With this revolt against the Fregean tradition comes the distinction between the context of employment of a token and the circumstance of its evaluation. As we saw, the content of a non-directly referential expression is a property, represented as a function from possible circumstances of evaluation to extensions. But the content of a directly referential expression, given the context, is just the actual referent no matter what the circumstances of evaluation are. Thus, in

(4) That is on a mat,

accompanied by a determination of a cat as the referent for 'that', the content of the verb phrase is the property of being on a mat, represented by a function from possible circumstances to classes of things on mats. But the content of the noun phrase, here a demonstrative pronoun, is the cat itself, which is the actual demonstratum given the context of the utterance. In this way the content of the whole becomes the singular proposition consisting of the cat

and the property of being on a mat. We may say somewhat figuratively that the expression ‘that’, being directly referential, is out of the game of being evaluated relative to non-actual circumstances of evaluation.

In the preceding two paragraphs and in all that follows I take certain liberties in terminology in order to adapt Kaplan’s framework to suit my own concerns. This deviation in terminology demands a word of explanation. Unlike Kaplan, my main interest here is not in the logic of demonstratives but rather in what may be regarded as their epistemology. Accordingly, I focus on acts of token-employment, normally utterances and inscriptions, rather than on what Kaplan calls ‘occurrences’, which are pairings of linguistic types and contexts of use in Kaplan’s sense. I regard contexts of use as possible contexts of token-employment, rather than as abstract bundles of parameters for determining the referents of occurrences as he does. And I regard characters as rules for token-employment, yielding contents in contexts of use in my sense, rather than as functions from contexts to contents. These points must be borne in mind for the remainder of the paper.

The above two-tiered framework for indexicals and demonstratives gives rise to a host of sentence-types that are of special interest to us, the so-called contingent analytic truths. Take the sentence ‘I am here’, where ‘here’ is an indexical. It is just in virtue of the characters involved that employing a token of ‘I am here’ cannot fail to express a true content. And insofar as it is the characters that the speaker can be said to know in knowing the language, it seems natural to say that it is in virtue of meanings alone that a token of ‘I am here’ is guaranteed to come out true in every context in which it is employed. On the other hand, it seems equally clear that it is not necessary for the speaker to be at a certain location at a certain time. She might well have been elsewhere at that time. In this way ‘I am here’ is said to be both analytic and contingent. Analyticity is a feature of the character associated with the type, whereas contingency is a feature of the content expressed by the token in context.¹¹

Now, in order not to prejudge the question whether contingent analytic truths may involve demonstratives, we characterize contingent analyticity in the following neutral way. A contingent analytic truth is a sentence-type, accompanied perhaps by a type of determination of referent for each of its demonstratives, such that: 1) contents expressed by tokens of this type are guaranteed to be true, and 2) the contents expressed by these tokens are not necessarily true. There are many examples of such types. What they all have in common is that while tokens of them express contents that have to be true in their respective contexts of use, what they express could have been otherwise. A historically important case in point to which I return in more detail below is Descartes’ ‘I exist’. Let the character of ‘I’ be given by the rule assigning the employer of the sentence-token containing the token of ‘I’ to the token of ‘I’ as its content. Then any employment of a token of ‘I exist’

is guaranteed to express a true content. But there are also clearly possible circumstances in which whomever ‘I’ refers to does not exist. Thus, even though ‘I exist’ is guaranteed by the characters involved to express a truth whatever the context of use may be, the parents of the referent of ‘I’ might never have met.

As mentioned above, demonstratives, as opposed to indexicals, require supplementation by a determination of the referent. Kaplan’s original account has this determination proceed via a demonstration, where a demonstration is a perceptual perspective on the demonstratum. On Kaplan’s later account such a determination is a directing perceptual intention rather than a demonstration. So, for example, with respect to (4), the demonstrative ‘that’ is now thought to be accompanied by the intention that a certain salient cat in the shared perceptual environment be taken as the referent. However we choose to think of a determination of the referent for a demonstrative, the kind of completion involved here can be represented by a description of the appearing of a purported demonstratum, a description that doesn’t contribute to the overall content of the token. Accordingly, we mark the supplementation of the demonstrative by such a determination of referent with the usual bracketing of the description of the appearing of the purported demonstratum. Following the demonstrative (say, ‘that’), a description of the appearing of the referent (‘the item appearing from here now as _’) will follow within brackets. The character of the demonstrative completed by such a description is then given by the rule that in any context, ‘that[the item appearing from here now as _]’ is a directly referential term designating whatever the description ‘the item appearing from here now as _’ picks out, if anything, and otherwise designating nothing.

Within this framework of demonstratives and determinations of referents, we may now consider whether there are contingent analytic truths involving demonstratives (D-CATs) beside those involving only indexicals. This requires us to consider sentence-types along with types of determination of referent, which are hybrids of the linguistic and the non-linguistic. A likely candidate may initially seem to be:

- (5) This[the item appearing from here now as the cat right in front of me] is here now.

But in a context in which there is really nothing in front of me despite it appearing to me that there is, the token in question would certainly not express a true content. A better candidate would be a sentence-type accompanied by a type of determination of referent involving self-reference. Let us turn to consider such cases.

As a preliminary, we need to inquire into some of the details of token-employment. This opens up an area of philosophical concern to whose depth I cannot do justice here. We need to ask: What must the relation

between a speaker and a token be in order for the speaker to succeed in expressing anything by the token? Schematically, to express oneself by a token is to carry out an intention to employ the token in some way. If this is correct, then a certain epistemic intimacy has to obtain between an employer of a token and the token employed in order for full-fledged expression to be achieved. Generally, if a token *appears* to me to be the one I am employing, then it should also *be* the one I am employing in order for me to succeed in expressing myself. Granted, even without actually looking at the words I am typing I can intend to employ the words on my screen.¹² But expressive success in cases where some token does appear to me to be one I am employing depends on the token so appearing being in fact employed by me. Otherwise, no saying is achieved after all. In other words, a speaker's expressive intention misfires if the appearance of an item employed in the expressive act is only a false appearance.¹³

Having gone through these considerations, we may now consider the following candidate for contingent analyticity:

- (6) This[the item appearing from here now as the sentence-token I am employing here and now] is here now.

Self-reference, i.e. the fact that the demonstrative is determined to refer to a token of which it is a part, makes it so that in each and every context in which a token of (6) is employed, the associated determination of the referent is bound to specify a demonstratum for the 'this' at the place to which 'here' refers at the time to which 'now' refers. Or consider the following example:

- (7) I am employing this[the item appearing from here now as the sentence-token I am employing].

Again, self-reference makes it so that in every context in which a token of (7) is employed, the associated determination of the referent is bound to specify as demonstratum a sentence-token employed by the referent of 'I'. So we may indeed take (6) and (7) to be D-CATs. Accordingly, the category of contingent analytic truths should be taken to include sentence-types accompanied by types of determination of referent.¹⁴ This possibility will prove essential for what follows.

4. Disquotation

We are now in a position to bring the demonstrative approach to quotation couched in the above framework of demonstratives to bear on the disquotation of, let us say, the name 'Socrates'. How does

(8) 'Socrates' refers to Socrates

fare under this approach? As a first approximation it should be analyzed as

(9) Anything sametokening that refers to Socrates. Socrates

where the second item is merely a demonstratum and thus has no linguistic role to play. But there is more to be said about the way in which (8) ought to be understood. As we saw earlier, it should be acknowledged that there is a sense in which instances of disquotation are guaranteed to come out true even if they do not express necessary truths. In order to see how this might be so I now propose to utilize the idea that demonstratives, as opposed to indexicals, require completion by a determination of referent.

Suppose that in thinking about philosophy I type a demonstrative paraphrase of (8) and my computer screen is so configured that the writing is in some bright color against a white background. I type a token of the first item of (9), the full sentence-token, and pause. Now suppose that you are reading over my shoulder so that your gaze becomes fixed on the last word-token. After a while, as you glance at the portion of the screen immediately following the sentence-token I had employed, you see an after-image so vivid that you are actually misled to believe that I had proceeded to type 'Socrates' there (despite the difference in color). Have I succeeded at this point in expressing a truth by what I had written? Answer: That depends on the associated determination of the referent. If the determination was to be of an item following the period, then I surely failed. There is just no token there to act as my demonstratum, even if you seem to see one. But if the determination was to be of the last word-token appearing on my screen then I was in fact successful.

Now, our pre-analyzed (8) is unlike its standard demonstrative paraphrase (9) in that in any context in which a token of it is employed the determination of the referent for the demonstrative is bound to have a demonstratum. It is just a feature of our orthographic conventions governing quotation that the quoted material is a proper part of the whole sentence-token, so that in every context in which the token of the whole sentence is employed the purported demonstratum is there to be referred to. Given that, we must now ask what determination of referent should be taken to accompany the demonstrative 'that' in an adequate demonstrative paraphrase of (8). And in light of the above after-image scenario, it is instructive to compare the following two options:

(10) Anything sametokening that[the item appearing from here now as immediately following the sentence-token I am employing] refers to Socrates. Socrates

(this is the analysis which is true to the letter of standard formulations of demonstrative paraphrases) and

- (11) Anything sametokening that[the item appearing from here now as the last word-token of the sentence-token I am employing] refers to Socrates.

I now claim that the sense in which instances of disquotation are guaranteed to come out true is generated by the specifics of the associated determination of the referent. The difference between (10) and (11) is not to be found at the level of the characters of the linguistic components. For these characters are in both cases the same. Nor is it to be found at the level of the contents of the tokens, provided that the demonstratum in each case is indeed an item such that anything sametokening it refers to Socrates. Rather, the difference in question is a difference in the character of the description of the appearing of the purported referent. And it is in virtue of this character, the character of the bracketed item following the ‘that’ representing the determination of the referent, that (8) is such that there is a sense in which it is guaranteed to come out true no matter what the context might be.

Understanding instances of the disquotational schema along the lines of (11) provides a perspicuous representation of the sense in which they are guaranteed to come out true even if they are not necessarily true. A fixed aspect of any context in which a token of (11) might be employed is that the last item of this token, which also provides the demonstratum, is being used to refer to whatever such items are being claimed to refer to. And as with (6) and (7) above, it is self-reference which secures the status of (11) as a D-CAT. In any context of use, the determination of the referent for the ‘that’ in (11) is guaranteed to have as demonstratum an item that is being used to refer to Socrates while being claimed (along with any such item) to refer to Socrates.¹⁵ But this does not make the contents of tokens of instances of disquotation necessarily true. Surely it is possible that the demonstrated item in our case would not have referred to Socrates, in which case items sametokening it would not either. For to be necessary, the content of a token of this instance of disquotation would have to be true in every possible circumstance. Yet there is nothing that necessitates that some particular demonstratum be employed in some particular way. The demonstrative theory of quotation is generally silent as to how the demonstrative gains its demonstratum. It is usually simply assumed that the quotation marks signal that some demonstrative reference is being achieved with the quoted material as demonstratum. Accordingly, demonstrative analyses of quotation generally place the specified demonstratum outside (either before or after) the item representing the logical form of the original locution. But if we adhere to this practice in the case of instances of

disquotation we may lose sight of the sense in which they are trivially and automatically true. By utilizing the above theory of demonstratives for this purpose we can actually represent the mechanism by which instances of disquotation are analytic while tokens of them do not express necessary truths.

Before turning to apply these results more widely, a clarification is in order concerning my construal of the particular way in which the demonstrative within a demonstrative paraphrase of disquotation gains its demonstratum. I have been urging that quotation is best understood via a demonstrative analysis and I have shown how instances of disquotation can be viewed as D-CATs by recasting them along the lines of (11). What we gain by doing so is perspicuity. But recasting them along such lines is not the only option available to us, even if we do choose to take a demonstrative approach to quotation and wish to view the likes of (8) as D-CATs. As mentioned above, the demonstrative approach imposes no restrictions on how the demonstrative is to gain its demonstratum as long as a demonstratum is there for the demonstrative to refer to. Accordingly, it would be possible for us to stipulate that the determination of the referent for the demonstrative within demonstrative paraphrases is uniformly of an item following the full sentence-token. We would then need to stipulate further that for the purpose of evaluating the content of a token of a demonstrative paraphrase in context, the two items of the paraphrase (the full sentence-token and the demonstratum) should be considered a single item. In that case, the fact that (8) is a D-CAT would be explained by the fact that the demonstrative in (10) is guaranteed to refer to a token that bears the sametokening relation to a different token that is being used to refer to Socrates.

What secures the status of (10) (and hence of (8)) as a D-CAT under this alternative construal is not, as in (6) and (7), self-reference. For under the new stipulation, according to which the demonstratum counts as a proper part of the sentence-token containing the demonstrative, there is a trivial sense in which every demonstrative paraphrase of quotation involves self-reference. What secures the status of (10) as a D-CAT in this case is the fact that (10) contains two tokens sametokening one another, one being used to refer to Socrates while the other being claimed (along with any other token sametokening it, including the first token) to refer to Socrates. What we seem to gain by choosing this sort of explanation for the contingent analyticity of (8) is surface uniformity among all demonstrative paraphrases of quotation in the character of the bracketed description. What we lose, however, is uniformity at the level of demonstrative reference. For now the demonstrative reference involved in demonstrative paraphrases is not just any old demonstrative reference. It is an intra-linguistic kind, with a built-in demonstratum. In the uniformity thus imposed on the determination of the referent, the demonstrative will become quasi-indexical. Yet if we can avoid introducing a special form of context-sensitivity to account for the workings

of quotation, we should. Moreover, by choosing this sort of explanation we will need to impose a certain restriction on the sametokening relation to the effect that the two tokens involved in the paraphrase will indeed have to sametoken one another. Weighing the alternatives against one another, I opt, then, for forgoing the new stipulations and construing instances of disquotation along the lines of (11). (8) is a D-CAT in virtue of the fact that (11) is, and (11) is such that a single token is both used to refer to Socrates and claimed, along with any item sametokening it, to be used in precisely that way.

I close this section by considering briefly a competing construal of instances of the disquotation of names as analytic. According to the metalinguistic theory of names, a name 'N' is semantically equivalent to a metalinguistic description such as 'The referent of "N"'.¹⁶ In this way, "'N'" refers to N' becomes semantically equivalent to "'N'" refers to the referent of "N'", which seems straightforwardly analytic. The theory seizes upon the triviality of instances of disquoting names as evidence against their direct referentiality. If an instance of disquoting a name is trivially true, how could the name lack a sense? One of the advantages of my account of disquotation is that it allows us to maintain that instances of disquotation are analytic without needing to give up the notion that names are directly referential.

A striking and highly disappointing feature of the metalinguistic theory is that even though it depends so heavily on quotation in its specification of metalinguistic senses, the semantic treatment of quotation itself is left completely unaccounted for. This has to be considered a severe theoretical drawback. By analyzing quotation along the lines of the demonstrative approach and by utilizing the above framework of demonstratives, we can both attain a better understanding of the ubiquitous device of quotation and offer an explanation for the triviality of the likes of (8) without needing to posit metalinguistic senses. And we can do all this while retaining the thesis that names are directly referential. In this way, sheer theoretical parsimony rules against the metalinguistic theory's explanation for the analyticity in question.¹⁷

But there is a further and to my mind decisive consideration against positing metalinguistic senses to explain the analyticity of instances of disquotation. Even though the metalinguistic theory would invoke a metalinguistic sense 'the referent of "Socrates"' to explain the trivial truth of (8), proponents of the theory do not make a similar appeal to a metalinguistic sense to explain the trivial truth of instances of disquotation involving common nouns, as in:

(12) 'cat' refers to cats.

The analyticity of such cases remains entirely unexplained.¹⁸ Yet the explanation for the trivial truth of (12) ought to be very much the same as the explanation for the trivial truth of (8).¹⁹

The approach to disquotation promoted here treats (12) and (8) on a par. (12) receives the following treatment:

- (13) Anything sametokening that[the item appearing from here now as the (singular form of the) last word-token of the sentence-token I am employing] refers to cats.

As with (11) in the case of disquoting the proper noun ‘Socrates’, it is self-reference that secures the status of (13) as contingent analytic. In any context, the determination of the referent for the ‘that’ in (13) is guaranteed to have as demonstratum an item that is being used to refer to cats while being claimed (along with any such item) to refer to cats. The uniformity of explanation for the analyticity of instances of disquotation involving proper and common nouns is just as we expect it to be. The metalinguistic theory’s failure to deliver such uniformity constitutes, I believe, a strong case against accepting its explanation for the analyticity in question.

5. Empty Names and Apriority

So far I have said nothing about disquoting empty names. If instances of disquotation are analytic in the manner suggested above, then their tokens would also seem to express a priori true contents. But in that case it would seem to follow that I am committed to claiming that a token of

- (14) ‘Vulcan’ refers to Vulcan

expresses a true content. Yet ‘Vulcan’ refers to nothing. To say that it refers shouldn’t be true, let alone true a priori. How should the discrepancy be resolved?

The answer here should be familiar from other discussions of empty names. Due to the emptiness of ‘Vulcan’, a token of (14) expresses no content. Accordingly, there is no knowing that ‘Vulcan’ refers to Vulcan, a priori or otherwise. Yet if I am right, a token of (14) is a token of an analytic type. What this shows is that tokens of analytic types may not express a priori truths. They can only express truths if they express contents to begin with.

It is useful to distinguish two different reasons for failure of apriority when it comes to knowing the truth of a content of a sentence-token. One reason is that no content at all is being expressed, in which case, in particular, one cannot know the content of the token to be true a priori. Consider the following analytic sentence-types:

- (15) It is not the case that both Vulcan has mass and does not have mass,

and

(16) Vulcan is self-identical.

Tokens of (15) and (16) would express contents knowable via the analyticity of the sentence-types—if only they expressed contents. Due to the emptiness of ‘Vulcan’, they happen not to. This is something that I believe we should learn to accept.²⁰ And it is one reason why a token of the analytic (14) fails to express an a priori truth. Such a token fails to express a content, true or false.

But there is another reason why a token of (14) fails to express an a priori truth. We just noted that tokens of analytic sentence-types may fail to express contents which are true a priori. But when it comes to sentence-types involving indexicals and demonstratives, there is an entirely different reason why tokens of them *must* so fail, a reason having nothing to do with the special case of empty names. Consider a token of the analytic sentence-type ‘I exist’ that I just employed. It would seem that by reflecting only on the characters involved I should be able to convince myself that this token expresses the particular true content of my existence. But this is surely too quick.

To simplify matters, suppose I say to myself, à la Descartes, ‘I exist’, by actually writing it down on a blackboard. Again, presumably I should be able to know a priori that the content expressed by this token is the truth of my existence by simply reflecting on the characters involved. But surely a random arrangement of molecules qualitatively identical to the token on my blackboard wouldn’t all by itself attest to my existence, nor will anyone else’s token of ‘I exist’. The rub lies in the conditions of employment of the token. It is simply not the case that by merely reflecting on the characters involved in ‘I exist’ I can convince myself of the particular truth of my existence. What should be said, rather, is that by reflecting on the characters involved I can convince myself that a token of ‘I exist’ expresses the truth of my existence, provided that it was employed by me in the right way. For me to come to know that the token of ‘I exist’ on my blackboard expresses the truth of my existence requires my knowing that it was employed by me in the right way. And this latter knowledge is not a piece of a priori knowledge. Specifically, it depends on my memory of the event of token-employment.

Once we see that knowing the specific content expressed by a token of ‘I exist’ depends on a posteriori knowledge, it is but a short step to the realization that this is not just a feature of the above ‘externalistic’ construal of token-employment. Even if I only employ a token of ‘I exist’ in my thoughts, I cannot conclude that I existed at the moment at which the token was employed without consulting my memory. The most I can conclude a priori is the conditional conclusion that *if* the token I am presently and inwardly demonstrating was employed by me in the right

way, *then* the content expressed by this token is true—I existed at the time of its employment. This is the most I can conclude. I cannot conclude that I existed then without relying on my memory of the event, which means that my existence then can only be a posteriori for me. And I surely cannot conclude that I exist now. These considerations point to an important qualification to traditionally inspired claims to a priori knowledge of one's own existence. In particular, they conflict with the familiar claim that one can know the truth of one's own existence a priori by running the cogito argument in one's own case. But this is a restriction on the achievement of the cogito argument of which Descartes himself was well aware. Seeing this merits a brief historical digression.

Descartes concludes his cogito argument with the following famous words: 'So after considering everything very thoroughly, I must finally conclude that this proposition, *I am, I exist*, is necessarily true whenever it is put forward by me or conceived in my mind' (Descartes 1988, 80). Insofar as Descartes' ruminations concern the type 'I am, I exist', he is no doubt correct and well ahead of his time in this famous observation of his. We would not choose to express the insight by saying 'necessarily true' and then qualifying it with the 'whenever' clause, but we see what he means. He is not attempting to advance the blasphemy that he is a necessary existent. He is attempting to show that his thought of his own existence is guaranteed to come out true. This is surely one of the most familiar conclusions in modern philosophy, but it is instructive to be reminded of how it is actually reached. Descartes offers the following:

I have convinced myself that there is absolutely nothing in the world, no sky, no earth, no minds, no bodies. Does it now follow that I too do not exist? No: if I convinced myself of something then I certainly existed. (Descartes 1988, 80)

Here is a rough reconstruction of the situation. He sits in his study and entertains the possibility that nothing is as it seems. He next reflects that if he actually managed to think that particular thought, then he had to have existed in order to think it, a conclusion expressed by: 'if I convinced myself of something then I certainly existed'. But notice the conditionality of this claim and notice the employment of the past-tense. The reason behind Descartes' double-caution here is not difficult to follow. First, with respect to the conditionality, at this stage of the *Meditations* we are already introduced to the possibility that memory itself might be systematically deceptive.²¹ So Descartes cannot simply detach the antecedent from the consequent and conclude that he existed at the time at which he authored his radical doubt because this depends on his remembering a certain event having taken place—his doubting. Second, with respect to the past-tense, Descartes must recognize that he cannot conclude by this argument that he exists at present. The most he can conclude is that if he had authored the

doubt that he is now inwardly demonstrating, then he had to have existed in order to author it. Reflection upon the thought that nothing is as it seems can only occur after its having been thought. There is always an ineliminable time-lag. First the thought occurs. Then the thinker reflects that for it to have occurred, he needs to have existed. Being in the past-tense, extracting from such reflection the conclusion that one did in fact exist then depends on one's memory of having authored the thought, that is to say, upon a posteriori knowledge of a past event.

We may summarize the point as follows. While 'I exist' is an analytic sentence-type, a token of it does not deliver the a priori knowledge of one's existence. The knowledge of one's existence through reflection on the semantics of 'I exist' depends on a posteriori knowledge that a given token of this type was employed in the right way by the very person who is reflecting upon it. Tokens do not wear their credentials on their sleeves. Yet it is precisely such credentials that are needed in order to interpret a token of an indexical or a demonstrative. So one conclusion to draw from all of this is that tokens of contingent analytic truths cannot deliver a priori knowledge. Coming to appreciate the contents of such tokens, i.e. interpreting them, inevitably involves relying on knowledge of the conditions of their employment. And this is a posteriori knowledge even as regarding tokens-in-thought.

Turning now to consider tokens of D-CATs such as (6) or (7), such tokens cannot express contents knowable a priori for a similar reason. Recall that D-CATs are hybrids of linguistic types and types of determination of referent, the latter being non-linguistic. Now consider the purely linguistic components of (6) and (7). In order to determine which content is being expressed by a token of 'This is here now' or 'I am employing this', we surely need to know what the 'this' in each case refers to. And this surely depends on a posteriori knowledge of the conditions of employment of 'this', just as it does in the case of 'I'. In other words, in order to identify a token of a given sentence-type as partaking in a token of a specific D-CAT—a hybrid of linguistic and non-linguistic elements—we need to know the conditions of employment of the linguistic token. And this cannot be a priori knowledge.

Returning to instances of disquotation involving empty names such as (14), we conclude that while such instances are analytic, tokens of them do not express contents at all, true or false, due to the emptiness of the name. This provides the first reason outlined above for refraining from taking such instances to be true a priori truths. But the second reason for contending that tokens of (14) could not express a priori true contents has nothing to do with the emptiness of 'Vulcan'. According to my proposal, tokens of the disquotational "'N" refers to N' are to be understood as tokens of 'Anything sametokening that refers to N', where 'that' is taken to refer to the last word-token. But when it comes to interpreting a token of 'Anything sametokening that refers to N', nothing betrays that the determination of

the referent for the ‘that’ is of the last word-token.²² Knowing that the content of a token of such a type is true requires knowing the conditions of employment of the demonstrative, and in particular, that the determination of the referent for the demonstrative is of the last word-token. And this surely depends on a posteriori knowledge. So tokens of instances of disquotation, just like tokens of ‘I exist’ and of ‘I am here now’, while tokens of analytic types, cannot deliver a priori knowledge.

6. Brains in a Vat

My construal of disquotation as contingent analytic has further epistemological ramifications for an improved understanding of Putnam’s BIV argument. Putnam’s nightmarish scenario is widely familiar, so I provide only a sketchy outline.²³ Suppose we are all, and have always been, brains in vats of nutrients whose nerve-endings are connected to a machine causing us to experience the collective hallucination of what we take to be our world. Putnam’s claim is that the skeptical hypothesis that we are brains in vats as described is self-refuting. His reasoning is as follows. Let me suppose that I am one of those brains. Then my words ‘brain’ and ‘vat’, say, would not refer to brains and vats, but rather to images generated by the machine, or perhaps to whatever features of the machine are causally responsible for the relevant images. But then ‘I am a brain in a vat’ as spoken by me will be false, because according to the generated image, from which the referents for the words in my language are drawn, I am not a brain in a vat (i.e. I do not appear to myself to be a brain in a vat). So under the supposition that I am a brain in a vat, ‘I am a brain in a vat’ as spoken by me is false. So I am not a brain in a vat.

It is commonly realized that what is doing the work in this argument is that one’s ability to describe the aforementioned scenario depends on one’s ability to refer to items which would be out of one’s referential reach had the scenario obtained. This realization gives rise to the following succinct version of the argument:²⁴

1. In my language, ‘brain’ refers to brains.
2. In vat-English, ‘brain’ does not refer to brains.

So,

3. my language is not vat-English.

And so,

4. I am not a brain in a vat.

An initial thought about this argument is that its first premise begs the question against the BIV skeptic. For even if we accept Putnam's claim that a BIV's 'brain' will not refer to brains, what is the status of the claim that in one's own language 'brain' refers to brains? Isn't this precisely what is at issue between the skeptic and the anti-skeptic? The skeptic would like to maintain that for all we know we might just be brains in vats, in which case our 'brain' would not refer to brains. But then we cannot just help ourselves to the truth of the first premise.

Having gone through some of the details of the demonstrative theory of quotation through the lens of the above framework of demonstratives, we may now assess the status of this first premise with more care. The first premise is true in virtue of the analyticity of a D-CAT. The type "'brain" refers to brains' is analytic, even though the fact that a particular token of 'brain' refers to brains is contingent. In terms of the account offered above, the analyticity here is a matter of the characters of the expressions partaking in a sentence-type involving a demonstrative together with the character of the bracketed description representing the associated determination of the referent, whereas the contingency is a matter of the content expressed by the token in context. A particular token of "'brain" refers to brains' expresses a true content in virtue of the fact that a word-token that is being demonstratively referred to, the likes of which are claimed to possess a certain semantic property, has this very semantic property. But for all that, a token of "'brain" refers to brains' does not express a necessary truth. Any word-token might have had a different semantic property than the one it actually has, or none at all.

In order to better assess the encounter between the BIV skeptic and the anti-skeptic, we need to remind ourselves of some of the background. A familiar theme in Putnam's work in the new theory of reference is that the rigidity of most common nouns derives from a certain indexicality in the fixing of their extension.²⁵ Take the familiar example of 'water'. Its extension is determined by the following condition:

(*) $\text{water}(x) \leftrightarrow \text{same}_L(x, \text{this})$.

The term 'water' applies to a given sample of a substance, in any world, just in case it bears a certain equivalence relation as befitting liquids to a paradigmatic sample referred to by an employment of 'this' in the actual world, say the substance in the glass on my kitchen table. So, claims Putnam, there is a hidden indexical component to 'water'. Water is whatever bears the sameness_L relation to the stuff around here, where sameness_L is roughly micro-structural identity. A parallel treatment applies to other common nouns. To be a brain is to bear a certain equivalence relation as befitting organs (sameness_O) to a paradigmatic sample referred to by an employment of 'this' in the actual world, say the organ steeped in formaldehyde in the jar on my kitchen table.²⁶

Going back to ‘“brain” refers to brains’ we now realize that we have a multiple context-sensitivity on our hands. First we have the following, including the demonstrative ‘that’:

- (17) Anything sametokening that[the item appearing from here now as the (singular form of the) last word-token of the sentence-token I am employing] refers to brains.

But if we next consider what ‘brain’ in (17) refers to, we should say, if Putnam is right, something along the following lines: It refers to anything bearing the sameness_O relation to a paradigmatic sample demonstratively referred to. Thus we have the following reference-fixing condition for ‘brain’:

- (**) brain(x) ↔ same_O(x,this).

Of course no one thinks that a proficient employment of ‘brain’ involves a conscious employment of (**) on the speaker’s part. The idea behind (**) is rather that when it comes to considering the general metasemantic question of how ‘brain’ refers (rigidly, as it happens) to whatever it refers to, something along the lines of (**) provides a schematic answer. In other words, ‘thing bearing the sameness_O relation to this[the item appearing from here now as the organ steeped in formaldehyde in the jar on my kitchen table]’ is not meant to *be* the content of ‘brain’, but rather to fix its content.

Given the above, the BIV scenario can now be brought into sharper focus. A token of ‘“brain” refers to brains’ as employed by me is guaranteed to express a true content in virtue of the analyticity of a certain hybrid of a sentence-type containing a demonstrative and a type of determination of referent for the demonstrative (a D-CAT). But *which* content such a token expresses is determined relative to the global context of my speech community. And what this global context is is not something that can be settled by my staring, however long, at a token of an instance of disquotation. To my mind, the interest in the BIV hypothesis lies in the way in which it capitalizes on the hidden indexicality of sentence-tokens that might not otherwise be suspected of involving any context-sensitivity, such as tokens of ‘“brain” refers to brains’. What my treatment of disquotation shows is that the global possibility spelled out in the BIV scenario cannot be ruled out on semantic and metasemantic grounds because such a possibility, had it obtained, would determine the very contents of the tokens in question. To see this, it is useful to consider an analogy.

David Christensen’s example of the desert man provides an instructive parallel to the workings of the BIV scenario.²⁷ You are standing alone in a vast featureless desert such that no feature of your current location could be evidentially distinguished from some location twenty feet to your right. You

then ponder the skeptical possibility that you are not really *here* now—perhaps you are really *there*? After all, the two locations are evidentially indistinguishable—were you there, things would seem to you exactly the same as they do right now.²⁸ Now, it may seem that by reflecting on the semantics of ‘I am here now’ one could put such a skeptical worry to rest. Surely you are located at the place to which ‘here’ refers, rather than at the place to which ‘there’ refers, at the time to which ‘now’ refers. But this is too quick a response to the worry.²⁹

You can know, by reflecting on the characters involved, that you could not employ a token of ‘I am here now’ to express a falsehood. But could such reflections settle where you are located in a way that would satisfy the location-skeptic? The answer is that they cannot. To know *which* content a token of ‘I am here now’ expresses requires knowing the credentials of the token. Specifically, it requires knowing where the token was located at the time of its employment. And this can only beg the question against the location-skeptic. In other words, an anti-skeptical strategy proceeding from the semantics of ‘I am here now’ would need to proceed in the following stages. First you would need to identify a given token as a token of the type ‘I am here now’ that you had just employed. Next, by reflecting on the type to which the token belongs, you should be able to conclude that the token in question indeed expresses a truth. Finally, in order to assess which true content this token expresses, you would need to know the location at which it was employed. In particular, you would need to know that you were not at that other location. And this is precisely what the location-skeptic is calling into doubt.

By analogy, an anti-skeptical strategy proceeding from the semantics of ‘“brain” refers to brains’ would have to proceed as follows. First you would need to identify a given token as a token of the type ‘“brain” refers to brains’ that you had just employed. Next, by reflecting on the type to which the token belongs and by the sort of considerations offered in this paper, you should be able to conclude that the token in question indeed expresses a truth. Finally, to assess which true content this token expresses you would need to know the global context of your speech community. In particular, you would need to know that the BIV scenario does not obtain.

In short, we can do justice to the BIV skeptic’s position to the extent that with respect to any token of ‘brain’ it is indeed not necessary that it should refer to anything in particular. And although the skeptic’s opponent is correct to point out that the type ‘“brain” refers to brains’ is an analytic truth very much akin to ‘I exist’ and ‘I am here now’, this cannot provide an effective response to the skeptic. Whatever we think about the achievement of semantic reflections in the cases of ‘I exist’ and ‘I am here now’ should carry over to this case as well. The construal of instances of disquotation as contingent analytic shows the limitations inherent in the semantic response to the BIV skeptic. It also shows the grain of truth in the conflicting

intuitions about the meaningfulness of linguistic expressions with which we began. That a token should have the semantic content that it does is surely a contingent fact about it. But it is a fact that can be appreciated via a token of a demonstrative analytic truth. As argued above, the analyticity here does not show that the meaningfulness of expressions can be appreciated in an a priori way. Having gone this far, the original intuition that it is a priori that an expression should have its content must give way to the notion that such a fact can be appreciated via the analyticity of a certain hybrid of a sentence-type containing a demonstrative and a type of determination of referent for the demonstrative. It is this analyticity that gives rise to the impression of apriority, and, arguably, it is not uncommon in philosophy for analyticity to be taken for apriority. I leave this larger issue for future work.

Notes

* For comments on and discussion of earlier versions of this paper I am grateful to Ido Geiger, Ruth Marcus, Michael Nelson, Shelly Rosenblum, Sanford Shieh, Abe Stone, and two anonymous referees. I am also indebted to participants in my Direct Reference seminar at Yale University during the Spring of 2001, and to audiences at the Second Bellingham Summer Philosophy Conference, Yale University, and UMass Amherst.

¹ I defer discussion of superficially similar cases involving empty names such as ‘“Vulcan” refers to Vulcan’ to section 5.

² I follow Kaplan’s distinction between semantics and metasemantics as outlined in (Kaplan 1989b, 573–576). Schematically, while semantics asks what the contents of linguistic expressions are, metasemantics asks how expressions gain their contents. For example, if the direct reference view of names is correct, proper names should be assigned their bearers as their contents. This is a semantic thesis. On the other hand, if the historical-chain view of names is correct, names gain their contents through historical chains linking current uses of names to initial episodes of naming. This is a metasemantic thesis.

³ Cf. (Sidelle 1991), (Williams 1993), and (Predelli 1998).

⁴ For the original statement of the theory see (Davidson 1979).

⁵ ‘On my theory, which we may call the *demonstrative theory* of quotation, the inscription inside does not refer to anything at all, nor is it part of any expression that does. Rather it is the quotation marks that do all the referring, and they help refer to a shape by pointing out something that has it’ (Davidson 1984, 90). The claim that quotations (expressions flanked by quotation marks) do not refer, a claim shared by all demonstrative approaches to quotation, has been the target of a recent criticism in (Gómez-Torrente 2001). The complaint is that because demonstrative approaches claim that quotations do not refer, they construe instances of higher-order disquotation such as

‘“Socrates”’ refers to ‘Socrates’

as false. This, it is claimed, conflicts with our ‘natural intuitions’ that such cases ought to come out true. But the alleged ‘intuitions’ are not genuine speakers’ intuitions at all. They are theoretically informed judgments about instances of high-flown philosophy-speak. Whether or not quotations refer is a theoretical matter, not one regarding which we should consult our intuitions. I plan to address this and related concerns about demonstrative approaches to quotation in greater detail in the future.

⁶Quine seems especially impressed with this use of quotation: ‘When we quote a man’s utterance directly we report it almost as we might a bird call’ (Quine 1960, 219).

⁷There are also general pragmatic reasons for thinking that such cases should not be accommodated by resorting to conversational implicatures. For discussions of some of the conditions for being included among the conversational implicatures of a given speech act, see (Carston 1988), (Recanati 1989), and (Recanati 1993).

⁸Cf. (Cappelen and LePore 1997), (Cappelen and LePore 1998), and (Cappelen and LePore 1999).

⁹(Simchen 1999). Cases of mixed quotation are cases in which the quoted material, besides being mentioned, also performs its ordinary task in the surrounding sentential context, as in: ‘Smith said that Jones “will leave”’. The phenomenon of mixed quotation is arguably the most prevalent use of quotation in ordinary discourse. Accordingly, any theory of quotation adequate to natural language has to meet at least the following desiderata: (a) it should accommodate the richness of diversity in our quoting practices, and (b) it should accommodate mixed quotation. Cappelen and LePore’s account is the only theory meeting these requirements.

¹⁰The fullest statement of Kaplan’s treatment of demonstratives is found in (Kaplan 1989a) and (Kaplan 1989b). For Perry’s version see (Perry 1993).

¹¹Calling the phenomenon of yielding a true content in every context of use ‘analyticity’ may strike some as perverse, to the extent that analyticity, as traditionally conceived, derives from a feature of content, whereas in this case it clearly does not. But I, along with many others, find it useful to retain the term for its heuristic value.

¹²This does not entail, however, that I thereby also intend to employ any string of characters on my screen (as when my fingers are off the intended keys), or that I intend to employ the very same words on any other screen (as when without my knowledge I am also producing tokens of the same types on some screen in an FBI office off-campus). The issues here are intricate and not very well-understood, but it seems that my expressive intention is generally quite specific, even when I am ‘not looking’.

¹³Think of this as an Austinian requirement. If in your attempt to name a ship it turns out that the bottle you thought you were smashing against the ship’s side was merely a false appearance, then in an important sense you failed to name the ship (assuming, of course, that naming the ship requires such a ceremony). Analogously, if it turns out that the utterance ‘I hereby name this ship the “Queen Elizabeth”’ was produced, unbeknownst to you, only by a bystander, perhaps an accomplished bilquist, then in an important sense you failed to name the ship even if you intended to name the ship ‘Queen Elizabeth’. (Whether or not the ship was consequently named is a different matter.) I am grateful to Abe Stone and to Ido Geiger for stimulating discussions of some of the intricacies surrounding token-employment. Neither of them should be held responsible for my conclusions.

¹⁴Once again, using ‘analyticity’ for the phenomenon in question may raise some eyebrows. (See note 11.) In the case of D-CATs, which are hybrids of sentence-types and types of determination of referent, ‘analyticity’ might seem even less apt than in the case of contingent analytic truths of the purely indexical variety. This is so to the extent that, unlike the case of contingent analytic truths, where tokens of indexicals gain their contents semantically, the phenomenon of D-CATs depends on tokens of demonstratives gaining their contents meta-semantically. (See note 2.) However, the traditional view of analyticity is not one that can facilitate the distinction between semantics and metasemantics in the first place. To see this, recall that on the traditional view of content, content uniformly determines reference, which is why there can be no room for reference-fixing that is not at the level of content. But such reference-fixing is precisely the domain of metasemantics. Thus, by the lights of the tradition, regarding D-CATs as analytic is neither more nor less problematic than regarding contingent analytic truths as analytic. Insofar as we find ‘analyticity’ useful for the latter, as many of us do, we can expect to find it useful for the former as well.

¹⁵Here I am assuming that sametokening is at least reflexive. It is surely symmetric. The more subtle question of transitivity need not concern us.

¹⁶ Proponents of various versions of this view have included Bach, Fodor, Katz, Kneale, Recanati, and Russell.

¹⁷ Notice also that if indeed the demonstrative theory of quotation is motivated independently of any considerations about the semantics of names—as I believe it is—and if Kaplan's two-tier framework for demonstratives is adopted, then *all* metalinguistic descriptions turn out to be synonymous! 'The referent of "N"' and 'The referent of "M"' have the same character, namely, the character of 'The referent of anything sametokening that'. The difference between the demonstrative paraphrases of the two descriptions is at the level of content.

¹⁸ Katz writes: '[T]he significant semantic difference between proper and common nouns is that the former have a sense that is metalinguistic' (Katz 1990, 5). For a similar rejection of metalinguistic senses for common nouns see (Recanati 1993, 161–165) and (Bach unpublished).

¹⁹ Curiously, Katz fails to acknowledge the trivial truth of instances of disquotation of common nouns. In (Katz 1990, 53) he considers whether or not 'horse named "Horse"' is redundant. It surely isn't. He ought to have considered whether or not 'horse that is a referent of "horse"' is redundant—it is.

²⁰ But I leave the pressing problem of negative existentials aside.

²¹ 'I will believe that my memory tells me lies, and that none of the things that it reports ever happened' (Descartes 1988, 80).

²² Similarly for a token of the pre-analyzed "'N" refers to N', it may be a typographical accident that whatever is enclosed within the quotes resembles the last word-token. My account of disquotation explains why this is no accident, but that it is no accident is precisely what we can't know a priori.

²³ It is described in full in (Putnam 1981, 1–21).

²⁴ This disquotational variant of the argument has been discussed at length by Crispin Wright. Cf. (Wright 1992, 67–94). I have chosen to simplify his rendition somewhat by suppressing some of his premises.

²⁵ The *locus classicus* of this view is (Putnam 1975a). Putnam's proposed analysis is meant to apply to all but a few hundred common nouns, the 'one-criterion words': nouns whose content is given by a simple definition (e.g. 'female fox' for 'vixen').

²⁶ I am leaving out of the present discussion any concern with the social aspect of the determination of content (i.e. the phenomenon of division of linguistic labor).

²⁷ The example is discussed in (Christensen 1993).

²⁸ An obvious implication of the scenario is that what engenders the problem for the desert man has little to do with the assumed featurelessness of the desert. After all, the man can locate himself where his nose is, but then again he won't be able to locate his nose. Similarly, he can locate himself where a certain rock is located without being able to locate the rock. In short, the desert man's worry is ultimately a worry about not being to locate himself in absolute space.

²⁹ Christensen actually takes the skeptical hypothesis in this case to be defeated by these semantic considerations: '[T]he referential devices [the desert man] is using work in such a way as to guarantee the truth of his claim "I am here"'. Were our man healthier, he could presumably have solved his skeptical problem easily, by adducing considerations of just this sort' (Christensen 1993, 311). As will become immediately clear, I disagree with Christensen's assessment.

References

- Almog, Joseph, Perry, John and Wettstein, Howard (eds.) (1989) *Themes from Kaplan*, Oxford: Oxford University Press.
- Bach, Kent. 'Descriptivism Distilled', unpublished manuscript.
- Boolos, George (ed.) (1990) *Meaning & Method: Essays in Honor of Hilary Putnam*, Cambridge: Cambridge University Press.

- Cappelen, Herman and LePore, Ernie. (1997) 'Varieties of Quotation', *Mind*, Vol. 106, 429–450.
- Cappelen, Herman and LePore, Ernie. (1998) 'Reply to Tsohatzidis', *Mind*, Vol. 107, 665–666.
- Cappelen, Herman and LePore, Ernie. (1999) 'Using, Mentioning and Quoting: A Reply to Saka', *Mind*, Vol. 108, 741–750.
- Carston, Robyn. (1988) 'Implicature, Explicature, and Truth-Theoretic Semantics', in (Davis 1991), 33–51.
- Christensen, David. (1993) 'Skeptical Problems, Semantical Solutions', *Philosophy and Phenomenological Research*, Vol. 53, 301–321.
- Davidson, Donald. (1979) 'Quotation', in (Davidson 1984), 79–92.
- Davidson, Donald. (1984) *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press.
- Davis, Steven (ed.) (1991) *Pragmatics: A Reader* (Oxford: Oxford University Press, 1991).
- Descartes, René. (1988) *Descartes: Selected Philosophical Writings*, Cambridge: Cambridge University Press.
- Gómez-Torrente, Mario. (2001) 'Quotation Revisited', *Philosophical Studies*, Vol. 102, 123–153.
- Harrison, Jonathan. (1998) 'The Trouble with Tarski', *Philosophical Quarterly*, Vol. 48, 1–22.
- Kaplan, David. (1989a) 'Demonstratives', in (Almog, Perry, and Wettstein 1989), 481–563.
- Kaplan, David. (1989b) 'Afterthoughts', in (Almog, Perry, and Wettstein 1989), 565–614.
- Katz, Jerrold. (1990) 'Has the Description Theory of Names been Refuted?', in (Bohoss 1990), 31–61.
- Kripke, Saul. (1980) *Naming and Necessity*, Cambridge: Harvard University Press.
- Perry, John. (1993) *The Problem of the Essential Indexical*, Oxford: Oxford University.
- Predelli, Stefano. (1998) '"I Am Not Here Now"', *Analysis*, Vol. 58, 107–115.
- Putnam, Hilary. (1975) *Mind, Language and Reality: Philosophical Papers*, Vol. 2, Cambridge: Cambridge University Press.
- Putnam, Hilary. (1975a) 'The Meaning of "Meaning"', in (Putnam 1975), 215–271.
- Putnam, Hilary. (1981) *Reason, Truth and History*, Cambridge: Cambridge University Press.
- Quine, W. V. (1960) *Word and Object*, Cambridge: MIT Press.
- Recanati, François. (1989) 'The Pragmatics of What Is Said', in (Davis 1991), 97–120.
- Recanati, François. (1993) *Direct Reference: from Language to Thought*, Oxford: Blackwell Publishers.
- Sidelle, Alan. (1991) 'The Answering Machine Paradox', *Canadian Journal of Philosophy*, Vol. 21, 525–539.
- Simchen, Ori. (1999) 'Quotational Mixing of Use and Mention', *Philosophical Quarterly*, Vol. 49, 325–336.
- van Fraassen, Bas C. (1997) 'Putnam's Paradox: Metaphysical Realism Revamped and Evaded', *Philosophical Perspectives*, Vol. 11, 17–42.
- Williams, C. J. F. (1993) 'Do I Have To Be Here Now?', *Ratio*, Vol. 6, 165–180.
- Wright, Crispin. (1992) 'On Putnam's Proof That We Are Not Brains-in-a-vat', *Proceedings of the Aristotelian Society*, Vol. 92, 67–94.