

Opacity and the attitudes

François Recanati

► **To cite this version:**

François Recanati. Opacity and the attitudes. Alex Orenstein, Petr Kotatko. Knowledge, Language and Logic: Questions for Quine, Kluwer Academic Publishers, pp.367-406, 2000. ijn_00000180

HAL Id: ijn_00000180

https://jeannicod.ccsd.cnrs.fr/ijn_00000180

Submitted on 14 Oct 2002

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

François Recanati
Opacity and the Attitudes*

Summary:

I. Opacity

1. Use vs. mention
2. 'Giorgione'
3. Referentiality and transparency
4. Transparency and substitutability

II. Belief sentences

1. Singular vs general beliefs
2. Scope ambiguities in belief sentences
3. Singular terms in belief sentences
4. The ambiguity of the *de re/de dicto* distinction

III. Opacity in belief sentences

1. 'That'-clauses as complex demonstratives
2. Semantic innocence
3. The attitudinatives as dependent expressions
4. Opacity, substitution, and quantification

Appendix

Footnotes

References

I. Opacity

§1.1 Use vs. mention

When we mention an expression, do we use it? It depends on *how* we mention the expression in question. We can mention an expression A by using another expression B which names it. In such a case we are not using A, but its name ('heteronymous mention'). But we can also use A itself in '*suppositio materialis*', that is, autonomously. That is what is ordinarily called 'mention' as opposed to 'use'. This traditional contrast between use and (autonomous) mention should not make us forget that in autonomous mention, the mentioned word itself is used, though deviantly.

I have just exploited the contrast between autonomous and heteronymous mention to lend credit to the idea that in autonomous mention (e.g. (1) below), the word is used. That is indeed the basis for the contrast; for the mentioned word does not occur at all in sentences such as (2) in which it is *heteronymously* mentioned.

- (1) 'Cat' is a three-letter word
- (2) Wychnevetsky is a three-letter wordⁱ

Still, it may be argued that the word 'cat' occurs only accidentally in (1), much as 'nine' occurs in 'canine' or 'cat' in 'cattle'. Such a claim, made by Quine, would be highly implausible if taken at face value.ⁱⁱ The occurrence of 'cat' in 'cattle' is indeed an accident; so much so that the word 'cat' does not, *qua* word, occur in 'cattle'. A sequence of letters (or a sound) is not a word. To be sure, the individuation of words raises complex issues, but on any plausible account 'cat' in 'cattle' will not count as an occurrence of a word. In contrast, the occurrence of 'cat' in (1) will count as an occurrence of the word 'cat', rather than as an orthographic accident. We can go along with Quine and accept that the first word of (1) is not the word 'cat', but a different

expression formed from the word 'cat' by appending quotation marks around it.ⁱⁱⁱ still, the occurrence of the word 'cat' *within* the complex expression is no accident. The word 'cat' is named by quoting *it*. That is how autonomous mention works. The mechanism of autonomous mention requires that we use *the word itself*, and put it within quotation marks. This is *toto mundo* different from a case of heteronymous mention (the word A is named by the word B) in which, by accident, B contains A in the manner in which 'cattle' contains 'cat'. (Thus instead of 'Wychnevetsky' in (2) we might have used another, no less arbitrary name, viz. 'Wychnecatsky', in which by accident the orthographic sequence 'cat' occurs. The difference with a standard case of autonomous mention is obvious.)

The mentioned word is used, but, as I said, it is used *deviantly*. The word is not used according to its normal semantic function. Thus a word whose role is to name a certain object or to "make it the subject of discourse" (as Mill says) will be used to make *itself* the subject of discourse.

Deviant uses, in general, are far from uncommon, and come in many varieties. We may not only use the word 'cat' autonomously, to denote that very word, but also to denote, say, a *representation* of a cat. Thus we can say:

(3) In the middle of the piazza stood a gigantic cat, due to a local sculptor.

This is deviant because a stone cat is not a cat. So the word 'cat', which means *cat*, can be used to mean many other things through the operation of various 'primary pragmatic processes' — pragmatic processes involved in the determination of what is said (Recanati 1993). Autonymy is one such process; metonymy is another. Such processes generate *systematic ambiguities*. Whenever a word denotes a type of thing, we can use it alternatively to denote a representation of that type of thing; whatever a word denotes, we can use it to denote that very word. Such ambiguities are similar to those mentioned by Quine himself (1960: 130): the process/product ambiguity (e.g. 'assignment' which can refer to the act of assigning or to the thing assigned), or the

action/custom ambiguity ('skater', which can refer to someone who skates or to someone who is skating), or the type/token ambiguity.^{iv}

In natural language, such ambiguities flourish. It is through them that natural language gains its main virtue: its flexibility, which makes it fit to talk about anything. But what is a virtue from one point of view is a defect from another. From a *logical* point of view, ambiguity is to be avoided. That means that, instead of using the same word to mean different things, we should use distinct words.

When it comes to quotation and autonomous use, the reform is easy. It proceeds in three steps:

Step 1: Whenever a word is used autonomously, make that explicit by using e.g. quotation marks around the word. (Even logicians did not respect that precept in the earlier part of this century. Frege and Quine were pioneers in this regard.)

Step 2: Consider the complex expression — word plus quotation marks — as a new word which names the original expression.

Step 3: Ignore the occurrence of the original expression, as if it were accidental; treat it as a fragment of the new expression, as 'cat' is a fragment of 'Wychnecatsky'.

I interpret Quine as urging this treatment of quotation as *desirable from the logical standpoint*, and as part of the 'reform' which has to take place before we can subject natural language sentences to logical appraisal. But that does not mean that the treatment in question is descriptively correct (or thought by Quine to be correct), as an account of the way natural language works. Natural language does not fear ambiguities, it rather welcomes them. In particular there is no doubt that it allows using a word to refer to that very word.

§1.2 'Giorgione'

Is the autonymous word referential or not? It depends in what sense. If we accept that the word refers to itself, then it is referential after all.^v Its referentiality can be checked using the Principle of Substitutivity. Replacing the autonymous word A, which refers to itself, by another, B, which also refers to A, preserves truth-value, as the possible transition from (1) to (2) shows. But the mentioned word is not referential in the *normal* sense: it does not refer to its normal referent. In what follows I will take 'referential' to mean just that: referential in the normal sense. An occurrence of a word is referential, in that sense, if and only if it refers to the normal referent of the word.

A term's being referential does not guarantee that the word can be replaced *salva veritate* by an occurrence of another word referring to the same object. For that to be guaranteed, Quine says, the occurrence at issue must be *purely referential* — the term must be used "purely to specify its object" (Quine 1960: 142). This qualification is necessary because Quine thinks there is a continuum of cases from pure non-referentiality, as illustrated by (1), to pure referentiality. Quine gives the following example:

(4) Giorgione was so-called because of his size

In such cases, Quine says, the word (here 'Giorgione') has a dual role. It is both mentioned *and* used to refer. It is a mixture of autonymy and referentiality. It is because the word 'Giorgione' is not used purely referentially that substitution of 'Barbarelli' for 'Giorgione' fails to preserve truth, despite the fact that Barbarelli and Giorgione are (were) one and the same person.

There is an apparent paradox in Quine's admitting such intermediate uses. For cases of autonymous mention are in principle eliminable in favour of heteronymous mention, in Quine's framework; but the occurrence of the word cannot be eliminated if, while mentioned, it keeps doing its normal referential work. Quine dispels the paradox

by construing intermediate cases as involving two occurrences consolidated into a single one: a purely referential occurrence and an autonymous (hence eliminable) occurrence. A perspicuous paraphrase makes the duality explicit. Thus (4) is rendered as

(5) Giorgione was called 'Giorgione' because of his size

I think Quine's insight that there is a continuum of cases between pure autonomy and pure referentiality is correct and important. (See the Appendix.) But his classification of the 'Giorgione' example in that category is misleading, for there is a sense in which the word 'Giorgione' in (4) is used purely referentially. To be sure, the word 'Giorgione' is mentioned in (4). But there is no inconsistency between holding that the word is used purely referentially, and holding that it is mentioned; for it is mentioned *heteronymously* in (4). Far from referring to itself, the word 'Giorgione' is referred to *by means of a different expression*, viz. the demonstrative adverb 'so' in 'so-called'. Hence the word 'Giorgione' itself is not used in two ways (referentially and autonymously); it is used purely referentially. *In contrast to autonymous mention, heteronymous mention is compatible with purely referential use.* This point can be driven home by splitting sentence (4) in two, as Kit Fine has suggested (1989: 253):

(8)

A: Giorgione was Italian.

B: Yes, and he was so-called because of his size

Who would deny that the occurrence of 'Giorgione' in A's statement is purely referential? The fact that B's statement contains an expression demonstratively referring to the name 'Giorgione' in no way conflicts with the purely referential character of the occurrence thus demonstrated.

Quine appeals to the failure of substitutivity as proof that the occurrence of 'Giorgione' in (4) is not purely referential. For if it were, it would be substitutable. Now, even though Giorgione is Barbarelli, substitution of 'Barbarelli' for 'Giorgione' does not preserve truth. Substitution of 'Barbarelli' for 'Giorgione' in (4) yields (9), which is false:

(9) Barbarelli was so-called because of his size

But this proof that the occurrence of 'Giorgione' in (4) is not purely referential rests on an equivocation. The fallacy of equivocation is presented as follows in Quine's *Methods of Logic*:

The two conjunctions:

(10) He went to Pawcatuck and I went along

(11) He went to Saugatuck but I did not go along

may both be true; yet if we represent them as of the form 'p&q' and 'r&¬q', as seems superficially to fit the case, we come out with an inconsistent combination 'p&q&r&¬q'. Actually of course the 'I went along' in (10) must be distinguished from the 'I went along' whose negation appears in (11); the one is 'I went along to Pawcatuck' and the other is 'I went along to Saugatuck'. When (10) and (11) are completed in this fashion they can no longer be represented as related in the manner of 'p&q' and 'r&¬q', but only in the manner of 'p&q' and 'r&¬s'; and the apparent inconsistency disappears. In general, *the trustworthiness of logical analysis and inference depends on our not giving one and the same expression different interpretations in the course of the reasoning. Violation of this principle was known traditionally as the fallacy of equivocation. (...)*

The fallacy of equivocation arises... when the interpretation of an ambiguous expression is influenced in varying ways by immediate contexts, as in (10) and (11), so that the expression undergoes changes of meaning within the limits of the argument. In such cases *we have to rephrase before proceeding.* (Quine 1962: 42-43; notation and emphasis mine)

By the same reasoning, it can be shown that the alleged failure of substitutivity exhibited by the occurrence of 'Giorgione' in (4) is merely apparent. Substitutivity fails, Quine says, because, although Giorgione was so-called because of his size, and $\text{Giorgione} = \text{Barbarelli}$, Barbarelli was *not* so-called because of his size. Paraphrasing Quine, however, we can respond as follows:

The two statements:

(4) Giorgione was so-called because of his size

(12) Barbarelli was not so-called because of his size

may both be true; yet if we represent them as of the form 'Fa' and ' \neg Fb', as seems superficially to fit the case, we come out with an inconsistency, since $a = b$. Actually of course the 'so-called' in (4) must be distinguished from the 'so-called' which appears in (12); the one is 'called *Giorgione*' and the other is 'called *Barbarelli*'.

When (4) and (12) are rephrased in this fashion they can no longer be represented as related in the manner of 'Fa' and ' \neg Fb', but only in the manner of 'Fa' and ' \neg Gb'; and the apparent inconsistency disappears.

What this shows is that the substitution of 'Barbarelli' for 'Giorgione' does preserve truth after all. The appearance that it does not is caused by the fact that "the interpretation of an ambiguous expression is influenced in varying ways by immediate contexts, ... so that the expression undergoes changes of meaning within the limits of the argument." If, following Quine's advice, we "rephrase before proceeding" we must substitute 'called *Giorgione*' for 'so-called' in (4) *before* testing for substitutivity; and of course, if we do so, we see that substitutivity does not fail. From (5) and the identity ' $\text{Giorgione} = \text{Barbarelli}$ ' we can legitimately infer (13):

(5) Giorgione was called 'Giorgione' because of his size

(13) Barbarelli was called 'Giorgione' because of his size

I conclude that 'Giorgione' in (4) is purely referential: substitution preserves truth, appearances notwithstanding. Yet the substitution which preserves truth is not any old substitution of coreferential singular terms, but substitution *under a uniform interpretation of whatever context-sensitive expression occurs elsewhere in the sentence*. This condition is crucial, for an apparent failure of substitutivity may be caused by the fact that the semantic value of some context-sensitive expression in the sentence changes as a result of the substitution itself. (That will be so in particular when, as in the 'Giorgione' example, the sentence contains an expression demonstratively referring to the singular term which undergoes substitution). When that is the case, the failure of substitutivity is consistent with pure referentiality. Only a failure of substitutivity under conditions of uniform interpretation provides a reasonable criterion of non-purely referential use.

In his discussions of opacity Quine does not adhere to his own policy of 'rephrasing before proceeding' when the sentence at issue is relevantly ambiguous or context-sensitive. Instead of using 'substitutivity' in the sense of 'substitutivity under conditions of uniformity', he uses it in the sense of 'substitutivity *tout court*'. In that sense the occurrence of 'Giorgione' in (4) is indeed not substitutable. I will hereafter follow Quine and use 'substitutable' in this way. My point concerning the 'Giorgione' example can therefore be rephrased as follows: Pure referentiality does not entail substitutability; hence failure of substitutivity cannot be retained as a criterion of non-purely referential occurrence.

§1.4 *Pure referentiality and transparency*

An occurrence of a singular term is *purely referential*, Quine says, just in case the term, on that occurrence, is used "purely to specify its object". In other words, the term's semantic contribution, on that occurrence, is its (normal) referent, and nothing else. To be sure, a singular term will not only contribute its semantic value (its referent), it will also show or display whatever other properties it has: its form, its sense, its affective

tone, its poetic qualities, and whatnot. But what matters from a semantic point of view is merely that which the term contributes to the truth-conditions of the whole.

What is meant exactly by a term's 'semantic contribution', i.e. its contribution to the truth-conditions? There is an ambiguity here. On a broad reading, the semantic contribution of an expression is *the overall difference it makes to the truth-conditions of the sentence where it occurs*. In that sense, 'Giorgione' in (4) does not make the same semantic contribution as 'Barbarelli' in (9); for if they did, (4) and (9) would have the same truth-value. But there is a stricter reading, more relevant to semantic theory. From the standpoint of semantic theory, each expression has a semantic value, and the semantic value of the sentence depends upon the semantic values of its parts and the way they are put together. The semantic contribution of an expression, in the narrow sense, is its semantic value — that which, in part, determines the truth-value of the whole. Thus in the 'Giorgione' example, what the word 'Giorgione' contributes is the individual Giorgione, which it names. The name 'Giorgione' serves also as referent for another expression, and affects the truth-conditions of the sentence in that respect too, but that is not part of the name's semantic contribution (in the narrow sense). Mentioning the name 'Giorgione' is something which *another* expression does; hence it is the semantic contribution of that other expression — while the semantic contribution of the name 'Giorgione' is the individual Giorgione, and nothing else.

As I am using it, the notion of a *purely referential occurrence* of a term is defined in terms of its narrow semantic contribution: a singular term is used purely referentially iff its semantic contribution is its referent, and nothing else. But there is room for a distinct notion, defined in terms of the 'broader' type of contribution. Let me define a *transparent* occurrence of a singular term as an occurrence such that *the semantic value of the sentence depends only upon the referent of the term, not on its other qualities (its form, its sense, etc.)*. Thus an occurrence is transparent iff its contribution in the broad sense is its referent, and nothing else.

The distinction between the 'broad' and the 'narrow' semantic contribution of a term, and correlatively between pure referentiality and transparency, is important

because it is possible for a term to be purely referential in a sentence, i.e. to contribute its referent and nothing else (in the narrow sense), without being transparent, i.e. such that the truth-value of the sentence does not depend upon any other quality of the term. For suppose that the sentence contains another singular term which demonstratively refers to the first one. Then, even if both terms are purely referential, the truth-value of the sentence will depend upon another property of the first term than merely its referent. That will not bar the first term from being purely referential since those aspects of the term, other than its referent, on which the truth-value of the sentence depends will not be part of the semantic contributions of *that term*, but part of the semantic contribution of the other term. That is exactly what happens in the 'Giorgione' example, as we have seen: though purely referential the term 'Giorgione' is not transparent; for the semantic value of the sentence depends not only upon the referent of the term, but also on its identity.

This analysis does not depend on my controversial construal of 'so' in 'so-called' as a demonstrative adverb. If we construe it as anaphorically linked to the name, the situation will be exactly the same: the semantic value of the sentence will depend upon the identity of the purely referential singular term *qua* antecedent of the anaphor. A striking example of that situation is provided by the following example, due to Kit Fine. He imagines a situation in which the man behind Fred is the man before Bill. Despite this identity we cannot infer (14) from (15):

- (14) The man behind Fred saw him leave
- (15) The man before Bill saw him leave

This does not show that the description 'the man behind Fred' is not used purely referentially; only that the occurrence of the description is not 'transparent', in the sense I have just defined.

To sum up, transparency entails pure referentiality, but not the other way round. There are *two* ways for an occurrence of a singular term not to be transparent.

- It can be non (purely) referential.
- The linguistic context in which the word occurs may be such that, even if it is purely referential, the truth-value of the sentence will depend upon other properties of the term than its referent. In this type of case I will say that the term occurs in a *reflecting context*; where a reflecting context is *a linguistic context containing an expression whose semantic value depends upon the identity of the term*.

In the second type of case, it's not the way the term is used but rather the context in which it is tokened that blocks substitutivity and generates opacity (the lack of transparency). Hence Quine's shift to talk of 'positions' instead of 'uses' or 'occurrences'. Quine defines a *position* as 'non purely referential' just in case the term in that position is not substitutable; this may be because the term itself is not being used in a purely referential manner, *or* because the linguistic context contains some context-sensitive expression whose value depends upon the identity of the singular term. Quine's notion of a non-purely referential *position* thus corresponds to my notion of an *opaque* occurrence. If I am right in my interpretation, Quine's talk of 'positions' was motivated by his realizing that opacity sometimes arises from the context rather than from the term itself. A term, in and of itself, may be as referential as is possible; if that term is demonstratively referred to by some other expression in the sentence, substitutivity will fail.^{vi}

§1.5 *Transparency and substitutability*

We have distinguished between a purely referential occurrence of a term, and a transparent occurrence (or, in Quine's terminology, an 'occurrence in purely referential position'). Now I want to consider a third notion: that of a *substitutable* occurrence of a singular term, that is, an occurrence of a singular term which can be replaced by an occurrence of a coreferential singular term *salva veritate*.

We have seen that a purely referential occurrence may fail the substitutivity test if it is not transparent (if the 'position' is not purely referential). At this point the question arises, whether we can equate substitutability and transparency.

The first thing we must note in this connection is that it is in fact possible for a purely referential term to be substitutable *without* being transparent. An example of that situation is provided by (16):

- (16) Cicero is the person commonly referred to by means of
the first word of this sentence.

There is no reason to deny that 'Cicero' is purely referential in this sentence. Its semantic value is the individual Cicero, which it names. But the sentence's semantic value results from the contributions of all constituents, including the demonstrative phrase 'this sentence'. Now the referent of the demonstrative phrase, hence the semantic value of the sentence, depends upon the identity of the singular term occurring at the beginning of the sentence. If you change the singular term, you change the sentence, hence you change the referent of the phrase 'this sentence', thereby possibly affecting the truth-value of the sentence. The singular term 'Cicero' is therefore not transparent, because the truth-value of the sentence depends upon the form of the name, even though its semantic contribution is nothing other than its referent. The form of the name affects the truth-value of the sentence *via* the semantic value of another singular term in the sentence. Despite this lack of transparency, the singular term is substitutable: if we replace 'Cicero' by 'Tully', we change the truth-conditions, but the truth-value does not change.

In a case like that, the singular term is substitutable for quite extrinsic reasons. Indeed it can be replaced by any other personal name *salva veritate*, whether that name is coreferential with 'Cicero' or not!

That a singular term can be substitutable without being transparent is not actually surprising. For a term can be substitutable without even being referential. Linsky (1967: 102) gives the following example:

(17) 'Cicero' is a designation for Cicero

In this sentence the first occurrence of 'Cicero' is (purely) autonymous, like the second occurrence of 'Giorgione' in (5). Yet it is substitutable: replacement of 'Cicero' by 'Tully' or any other name of Cicero in (17) is truth-preserving.^{vii}

Let us grant that transparency cannot be equated with substitutability. Can we at least maintain, following Quine, that transparency *entails* substitutability? It seems that we should. Paraphrasing Quine (1960: 242), we can argue that

If an occurrence of a singular term in a true sentence is transparent, i.e. such that the truth-value of the sentence depends only upon the object which the term specifies, then certainly the sentence will stay true when any other singular term is substituted that designates the same object.

Yet even that has been (rightly) disputed. What I have in mind is Kaplan's insightful discussion of what he calls "Quine's alleged theorem" in 'Opacity' (Kaplan 1986).

Kaplan argues that, technically, substitutability does not follow from transparency. But the same point can be made in a non-technical framework, by appealing to the same sort of observation which enabled us to draw a distinction between pure referentiality and transparency.

The crucial point, again, is that natural language sentences are context-sensitive to such a degree, that substituting a singular term for another one can affect the interpretation of other expressions in the same sentence. This may block substitutivity and generate opacity even if the terms at issue are purely referential. Now when a singular term is not only purely referential but *transparent*, it seems that no such thing can happen: for the context is (by definition) not reflecting; it does not contain

expressions whose semantic values depend upon the identity of the term. How then can the substitution of coreferentials affect the interpretation of the rest of the sentence? It seems that it cannot, yet, I will argue, it can.

Let us imagine a purely referential occurrence of a term t in a sentence $S(t)$, and let us assume that that occurrence is transparent in the sense that the truth-value of $S(t)$ depends upon the referent of t but not on any other property of t . Since the occurrence of t is transparent, the context $S()$ is not reflecting. Since it is not reflecting, it seems that if we replace t by a coreferential term t' , and if the occurrence of t' also is purely referential, then t' can only be transparent. The truth-value of $S(t')$ will therefore depend upon the referent of t' but not on any other property of t' . It follows that $S(t')$ will have the same semantic value as $S(t)$: t , therefore, is substitutable in $S(t)$.

But there is a hidden assumption in the above argument, an assumption which is in fact questionable. It is this: that the linguistic context $S()$ is 'stable' in the sense that if it is non-reflecting in $S(t)$, then it is also non-reflecting in $S(t')$. But suppose we lift that assumption; suppose we accept *unstable contexts*, that is, contexts whose interpretation can shift from non-reflecting to reflecting, depending on which singular term occurs in that context. Then we see that a transparent singular term may not be substitutable after all.

Let us, again, assume that the occurrence of t in $S(t)$ is transparent. This entails that, on that occurrence, t is purely referential and $S()$ is non-reflecting. Yet we cannot conclude that $S()$ will remain non-reflecting after we have substituted t' for t . For an unstable context is a context which is ambiguous between a reflecting and a non-reflecting interpretation. If $S()$ is unstable in this way, then it may be that $S()$ is non-reflecting in $S(t)$ but becomes reflecting in $S(t')$. Suppose that is the case; then t' is not transparent in $S(t')$: the truth-value of $S(t')$ will not depend merely upon the referent of t' — it will depend on the identity of the term. The truth-conditions, hence possibly the truth-value, of $S(t')$ will therefore be different from the truth-conditions of $S(t)$. In such a case, therefore, t is not substitutable: replacing it by a purely referential occurrence of a coreferential term t' may result in a change of truth-value!

That is not a purely theoretical possibility. There are reasons to believe that attitude contexts are unstable. A belief sentence like 'John believes that Cicero is bald' has two readings: a purely relational reading in which it says of John and Cicero that the former believes the latter to be bald, without specifying how (under which 'mode of presentation') John thinks of Cicero; and a non-purely relational reading in which it is further understood that John thinks of Cicero as 'Cicero'. According to several authors, who use the 'Giorgione' example as paradigm, 'John believes that ... is bald' is a reflecting context on the non-purely relational reading;^{viii} that is, the sentence somehow involves a 'logophoric' or demonstrative reference to the singular term which occurs in the context. Even if the term in question is construed as purely referential, the truth-value of the sentence depends not only on the referent of the term but also on its identity, on the non-purely relational reading. In contrast, the context is non-reflecting on the purely relational reading. If I say

John, who confuses me with my grandfather Frank Recanati, believes
that I died twenty years ago

the truth-value of the sentence depends only upon the referent of 'I'.

If belief contexts are ambiguous and unstable in this manner, which particular singular term occurs in the sentence may affect its interpretation. This blocks substitutivity: even if the occurrence of the singular term *t* in 'John believes that *t* is *F*' is not only purely referential but also transparent, substituting a purely referential occurrence of a coreferential singular term *t'* for *t* may shift the interpretation of 'John believes that... is *F*' to its reflecting reading, thereby making the occurrence of *t'* opaque. That is what apparently happens if we replace 'I' by 'François Recanati' in (18):

(18) John believes that I died twenty years ago

(19) John believes that François Recanati died twenty years ago

In both cases John is said to have a belief concerning François Recanati, to the effect that he died twenty years ago; but in the second case there arguably is a logophoric or demonstrative reference to the singular term. (19) can be paraphrased as:

(19*) John *so-believes* that François Recanati
died twenty years ago

That interpretation of the ambiguous 'believes' is natural when the singular term is the proper name 'François Recanati', while the pronoun 'I' rules out this interpretation for pragmatic reasons (McKay 1981; Recanati 1993: 399-401).

I am not presently defending this analysis of belief sentences; I will do so in the third part of this paper. That brief anticipation was only meant to illustrate the notion of an unstable context, that is, a context ambiguous between a reflecting and a non-reflecting reading. In the same way in which a purely referential occurrence may not be transparent if it occurs in a reflecting context, a transparent occurrence may not be substitutable if it occurs in an unstable context. Thus in (18) the singular term 'I' is not substitutable even though it is transparent, because the context is unstable.

To be sure, if, following Quine's general methodological recommendations, we get rid of context-sensitivity by suitably rephrasing the sentences we subject to logical treatment, then we automatically get rid of both reflecting and unstable contexts. It then becomes possible to equate (as Quine does) pure referentiality, transparency and substitutability. But, as we saw, Quine himself does not follow his own recommendations: he treats 'Giorgione' as non-purely referential and non-substitutable in (4), something which is possible only if we take the context-sensitive sentences "as they come" (Quine 1960: 158), without prior rephrasing. It is this policy which enables him to put in the same basket non referential (autonomous) occurrences of terms and referential occurrences in reflecting contexts. I have shown that if we take this line, then we should draw a principled distinction between pure referentiality, transparency, and substitutability.

II. Belief sentences

§2.1 *Singular and general beliefs*

In his classic paper 'Quantifiers and propositional attitudes' (1956), Quine made a distinction between "two senses of believing", as he then put it: the notional and the relational sense. That is both a distinction between two readings of belief sentences, and a distinction between two types of belief. The distinction is very intuitive, but it faces difficulties. In later writings Quine expressed skepticism toward the distinction, and more or less gave it up (Quine 1977: 10). Contrary to Quine I think the distinction can be saved. What follows is my reconstruction of it.

Let us start with the distinction between two types of belief. Some beliefs are purely general, others are singular and involve particular objects. As an example of a general belief, we have the belief that there are spies, or the belief that all swans are black. As Frege put it, those beliefs are about concepts, if they are about anything at all: the first is the belief that the concept 'spy' is satisfied by at least one object, the second is the belief that whatever satisfies the concept 'swan' satisfies the concept 'black'. But the belief that Quine was a student of Carnap is a belief about two individual objects: Quine and Carnap. Of this belief we can say: There is an x and there is a y such that the belief is true iff x was a student of y . We cannot say anything similar concerning the belief that there are spies: there is no individual object x such that that belief is true iff x satisfies a given predicate.

A singular belief is relational in the sense that the believer believes something *about* some individual. The relation of 'believing about' descends from more basic, informational relations such as the relations of perceiving, of remembering or of hearing about. All those relations are genuine *relations*. If John perceives, remembers, or hears about the table, there is something which he sees, remembers or hears about. Similarly, if John believes something about Peter, there is someone his belief is about.

Singular belief is based on, or grounded in, the basic informational relations from which it inherits its relational character. To have a thought about a particular object, one must be *'en rapport with'* the thing through perception, memory or communication. Pure thinking does not suffice. Thus inferring that there is a shortest spy does not put one in a position to entertain a singular belief about the shortest spy, in the relevant sense.

In terms of this distinction between singular and general beliefs, well-documented and elaborated in the philosophy of mind (see e.g. Evans 1982), I suggest that we define a relational belief report as one that reports the having of a singular belief; and a notional belief report as one that reports the having of a general belief.

How do we know whether a given sentence reports a singular or a general belief? Can we tell from the form of the sentence, or is each belief sentence ambiguous between the two readings? Before dealing with this important question (§2.2-3), we must pause to consider Quine's likely attitude toward the distinction between singular and general belief.

As my examples reveal, singular beliefs are typically expressed by means of singular statements such as 'Quine was a student of Carnap'; and general beliefs by means of quantified statements such as 'There are spies' or 'Every swan is black'. But Quine notoriously downplays the difference between the two types of statement. Singular statements, he holds, can be rephrased as general statements (Quine 1960: 178ff). Thus 'Cicero is bald' says no more and no less than: 'There is an x such that x is Cicero and x is bald'. The difference between the two statements is purely rhetorical, Quine says.

Quine's elimination of singular terms in favour of general terms is not intended as a wholesale elimination of singularity, however, but as a displacement of it. If there is some distinguishing feature which singular statements possess, that feature will automatically be transmitted to the 'general' statements into which singular statements are rephrased in canonical notation. Quine insists that nothing is lost in the manoeuvre — the elimination of singular terms concerns only superficial grammar:

It is felt... that the names differ from the predicates in their connotation of uniqueness, though predicates may just happen to apply uniquely. It is felt also that proper names lack connotation while predicates connote. Now these are traits of names that I simply transfer to the predicates, however unaccustomed the new setting. This is why I spoke of *reparsing*: the names can keep all their old traits except grammatical position. (Quine 1980: 173)

The difference between 'Cicero is bald' and 'There are spies', then, is not essentially structural (both, according to Quine, are best seen as quantified statements), but lies in the nature of the predicates involved: the first but not the second type of statement involves what we might call a 'singular predicate', viz. 'is Cicero'.

Corresponding to the original distinction between singular and general statements, we now have a distinction between singular and general predicates. Singular predicates are those predicates which inherit, or otherwise possess, the distinguishing features of singular terms. If, as I have suggested, relationality is the distinguishing feature of genuine singular terms, then singular predicates will possess that feature as well. That means that one cannot believe that the predicate 'is Cicero' is instantiated without being suitably related to (*en rapport* with) Cicero. From this point of view, the singular predicate 'is Cicero' is very different from a truly general predicate like 'is a spy' or 'is called *Cicero*'.

Earlier I used the following criterion of singularity:

Criterion C:

A belief (or a statement) is singular iff:

There is an x such that the belief (or the statement) is true iff ...x...

Can we still use that criterion, in Quine's framework? I think so. If we rephrase 'Cicero is bald' as

- (1) There is an x such that x is Cicero and x is bald

that is still singular by criterion C. For *there is a y such that* (1) is true iff there is an x such that $x = y$ and x is bald. The same thing cannot be said of a fully general statement such as 'There are spies'.

Quine would certainly object to the recurring of the proper name 'Cicero' within the predicate 'is Cicero', however. In order to complete the elimination of singular terms, the singular predicate must be construed as 'notationally atomic' (Quine 1980: 173). Still, Quine says, the predicate will inherit the traits of the eliminated name, including — presumably — its essentially relational character. If that is so, then we can still use criterion C. We can define a singular predicate as follows:

A predicate F is singular iff:
there is a y such that the belief that there is an x such that x is F is true iff ...y...

Any belief to the effect that such a predicate is instantiated will count as singular by virtue of criterion C.

I conclude that Quine's elimination of singular terms other than variables does not threaten the distinction between singular and general beliefs (or between singular and general statements). Even if it did, however, we would not be forced to choose between the singular/general distinction and Quine's policy of letting only variables refer. For there is an alternative to Quine's way of eliminating singular terms other than variables — an alternative which, far from undermining the singular/general distinction, captures it in a rather elegant and straightforward manner.

Part of Quine's motivation for his regimentation is his belief that "names, like predicates, serve to characterize the thing referred to" (1980: 172). Thus when I refer to Cicero as 'Cicero', I characterize him as (being) Cicero. That claim is not very convincing to someone who holds that proper names are non-connotative. Be that as it may, Quine himself accepts that a pronoun such as 'he' does not do much by way of characterizing its referent. The pronoun, he says, is "purely referential [and] utterly uninformative"; it "connotes only the sex and scarcely that" (1980: 165). So there is a

certain convergence between Quine and theorists of singular reference as far as pronouns such as 'he' are concerned: both parties accept that pronouns are vehicles of pure reference. Quine accepts that because he sees pronouns as the natural language counterparts of variables (1974: 93-101); the theorist of singular reference because he sees *demonstrative* pronouns as paradigmatic singular terms (Kaplan 1989). I think this convergence can be exploited to eliminate singular terms other than variables in a way which is more congenial to the theorist of singular reference.

A pronoun, in general, is very much like a variable. Some pronouns are like bound variables: 'If a man catches a lion, *he* is brave'. A demonstrative pronoun is more like a *free* variable, under a contextual assignment of value. The suggestion, then, is this. When we say 'He is brave', pointing to the hunter, the sentence which we utter is neither true nor false: it is an open sentence. By asserting it, however, we present it as *true of* the object we are demonstrating. The assertion is true *tout court* iff that is indeed the case, that is, iff the demonstrated object satisfies 'x is brave'. In such a case *there is an x such that our assertion is true iff x is bald*. On the other hand, when the sentence is a closed sentence such as 'There are spies', there is no x such that the assertion is true iff ...x...

This treatment of demonstrative pronouns can be extended to all genuine singular terms, including proper names. A proper name such as 'Cicero' can also be considered as a free variable. I refer the reader to Dever (forthcoming) for an elaboration of this view. In that framework, I think, the elimination of singular terms is conducted in a way that enhances the distinction between singular and general statements.

§2.2 *Scope ambiguities in attitude contexts*

How do we tell whether a given belief sentence reports the having of a general belief or the having of a singular belief? Quine thinks that a standard belief sentence like 'Ralph believes that Ortcutt is a spy' is ambiguous between the relational and the notional

reading; and that we can force the relational reading by 'exporting' the singular term: 'Ralph believes *of Ortcutt* that he is a spy'. When exportation is thus possible, existential generalisation is also possible: if Ralph believes that Ortcutt is a spy, in the 'exportable' sense (that is, if he believes of Ortcutt that he is a spy), then there is someone Ralph's belief is about.

Even though Quine's claim concerning the ambiguity of belief sentences between the relational and the notional reading has been very popular, I think that it rests, in part, on a confusion; a confusion which is, again in part, responsible for Quine's despair of the distinction. In the next section I will argue that standard belief sentences such as 'x Vs that p', where the embedded sentence contains a singular term, are not ambiguous between the relational and the notional reading. That ambiguity arises only when the embedded sentence contains a quantified or descriptive phrase.

The distinction between genuine singular terms and descriptive or quantified phrases such as 'some man', 'a man', 'no man' or 'the man' goes back to Russell (1905). While Russell wanted to restrict the class of 'logically proper names' (as he called genuine singular terms) to only a couple of natural language devices, contemporary semanticists consider ordinary proper names and demonstratives, in general, as genuine singular terms. *Qua* genuine singular terms, they are purely referential, in the sense of §1.3. Definite and even indefinite descriptions can also be *used* purely referentially, according to some authors at least (Donnellan 1966; Chastain 1975); but the purely referential use of descriptions is not their normal semantic function, while it is the normal semantic function of genuine singular term.

There is a good deal of controversy over the referential use of definite descriptions. Many people believe that it is irrelevant to semantics. I disagree, but we need not be concerned with this issue here. If, as I believe, definite descriptions have a non-deviant referential use,^{ix} then, when so used, they behave like genuine singular terms: they are purely referential and their semantic value (on that use) is their referent. What I have to say about the behaviour of genuine singular terms in belief contexts will therefore automatically apply to definite descriptions on their referential use. So I will

put referential descriptions aside and consider only what Evans called the 'pure' uses of definite descriptions, that is, their attributive uses.

As Russell pointed out in the above-mentioned paper, definite descriptions are very much like quantified phrases. Like them, they serve to make general statements. If John believes or asserts 'The winner will be rich', we cannot say that there is an object x such that John's belief or statement is true iff x satisfies F , whichever predicate we put in place of the schematic letter ' F '. In particular, we cannot say that a certain person, namely the winner, is such that John's belief is true iff *she* will be rich; the condition 'being the winner' must also be satisfied by her. Nor can we say that a certain person is such that the belief is true iff she is both rich and the winner. *Any* person's being rich and the winner can make the belief true.

Definite descriptions are similar to quantified phrases in another respect: like them, they induce scope ambiguities in complex sentences containing an intensional operator. Thus there are two readings for sentences such as (2) or (3):

- (2) Someone will be in danger
- (3) The President will be in danger

(2) says either that someone is such that she will be in danger, or that it will be the case that someone is in danger. The two readings can be represented as follows:

- (2a) $(\exists x)$ (it will be the case that $(x$ is in danger))
- (2b) It will be the case that $((\exists x)$ $(x$ is in danger))

The same duality of readings can be discerned in the case of (3). (3) says either that the President is such that he will be in danger, or that it will be the case that: the President is in danger. On the second reading it is the fate of a future president which is at issue, while on the first reading the sentence concerns the present president. Again, the two readings can be represented in terms of relative scope:

- (3a) $(\ulcorner \text{President } x \urcorner) (\text{it will be the case that } (x \text{ is in danger}))$
 (3b) It will be the case that $((\ulcorner \text{President } x \urcorner) (x \text{ is in danger}))$

In (2a) and (3a), the quantifier or descriptive phrase is given wide scope; thus it seems to reach into the intensional context created by the operator 'it will be the case that'. But, as Kaplan (1968, 1986) and Quine (1977) pointed out, (3a) and (2a) need not be construed as actually violating Quine's prohibition of quantification into intensional contexts. The intensional operator 'it will be the case that', or 'will-be' for short, can be thought of as *multigrade* (Quine 1956, 1977). Its taking narrow scope vis-a-vis the descriptive or quantified phrase in (2a) and (3a) means that it governs only the predicate 'in danger', while it governs the whole sentence 'someone is in danger' or 'the President is in danger' when it is given wide scope, as in (2b) and (3b). That can be made notationally explicit in the manner of Quine 1977:

- (2a') $(\exists x) (\text{will-be}(\text{in-danger } x))$
 (2b') Will-be $((\exists x) (\text{in-danger } x))$
 (3a') $(\ulcorner \text{President } x \urcorner) (\text{will-be}(\text{in-danger } x))$
 (3b') Will-be $((\ulcorner \text{President } x \urcorner) (\text{in-danger } x))$

In (2a') and (3a') the multigrade 'will-be' is understood as a predicate functor making a new predicate, 'will be in danger', out of the original predicate 'in danger'. The quantified variable thus falls outside the scope of the intensional operator. When the operator is given wide scope, as in (2b') and (3b'), it is understood as governing the whole sentence (including the quantifier and the variable). The quantified variable now falls within the scope of the operator, but, as Quine says, the sentence "exhibits only a quantification *within* the 'believes that' context, not a quantification *into* it" (1956: 188).

Before proceeding, let us note that genuine singular terms give rise to no such scope ambiguities: they are, as Geach once put it, "essentially scopeless" (Geach 1972: 117). Thus sentence (4) is not ambiguous, contrary to (2) or (3); there is no truth-conditional difference between (4a) and (4b), as there was between (2a) and (2b) or between (3a) and (3b):

- (4) Cicero will be in danger
- (4a) Will-be (in-danger Cicero)
- (4b) Will-be(in-danger) Cicero

It is time to introduce belief sentences. Belief sentences with descriptive or quantified phrases are ambiguous in a way that exactly parallels the ambiguities we have just observed in temporal sentences with descriptive or quantified phrases. Thus (5) is ambiguous like (2), and (6) is ambiguous like (3):

- (5) John believes that someone is F
- (5a) Someone is such that John believes him to be F
= $(\exists x) (B_j(F) x)$
- (5b) John believes that: someone is F
= $B_j ((\exists x) (Fx))$

- (6) John believes that the President is in danger
- (6a) The President is such that John believes him to be in danger
= $(\exists x \text{ President } x) (B_j(\text{in-danger } x))$
- (6b) John believes that: the President is in danger
= $B_j ((\exists x \text{ President } x) (\text{in-danger } x))$

The quantification is endorsed by the speaker in (5a), while it is ascribed to the believer in (5b). Similarly, the description is endorsed by the speaker in (6a), while it is ascribed to the believer in (6b).

Note that in (3a) the description is to be read attributively even though it takes wide scope (Kripke 1977: 258). The speaker says that the G, whoever she is, is such that John believes her to be F. The description does not behave like a singular term here; it does not contribute an object. Still the *ascribed* belief is singular: the speaker says that there is a particular object such that the believer believes something of that object.

To sum up, when the quantified phrase or the description takes wide scope, belief reports like (5) and (6) have their relational reading: the belief they report is singular, even though the object the belief is about is only described in general terms.^x In contrast, when the descriptive or quantified phrase takes narrow scope, the belief report is understood notionally. The believer is said to believe that there is an object x with such and such properties; that does not entail that there actually is an object y such that the believer believes that of y. Whatever quantification there is is strictly internal to the ascribed content; it is not endorsed by the speaker.

§2.3 *Singular terms in belief sentences*

So far, Quine's claim concerning the ambiguity of belief sentences has been vindicated. But quantified phrases and definite descriptions are not genuine singular terms (Neale 1990). As soon as what occurs in the embedded sentence is a genuine singular term (or a referential description), the scope ambiguity vanishes, along with the distinction between the notional and the relational reading of the belief report.

Since a singular term is purely referential (unless it is used deviantly), a statement in which it occurs is bound to be singular. That is true not only of a simple statement such as 'Cicero is in danger', but also of a complex statement such as 'John believes that Cicero is in danger'. The former is about the individual Cicero; the latter is

about two individuals, John and Cicero. It follows that exportation is always licensed when the embedded sentence contains a genuine singular term.^{xi} From:

(7) John believes that t is F

we can always go to

(8) John believes of t that it is F

and, through existential generalisation, to

(9) $(\exists x) (B_j(F)x)$

That means that the ascribed belief is always singular, when the belief report contains a singular term. 'Notional' readings are thus ruled out: only relational readings are available.

What I have just said, of course, presupposes that genuine singular terms are used normally (non-deviantly) in attitude contexts. That is, I am assuming what Davidson (1968) and Barwise and Perry (1981) call 'semantic innocence' (see §3.2 below); and correlatively rejecting the notion that singular terms in attitude contexts refer to something different from their usual referent (Frege) or behave somewhat deviantly, as they do when they occur autonomously (Quine). I take singular terms to be purely referential, in the sense of §1.3, in all their non-deviant occurrences; and I assume that their occurrences in attitude contexts are non-deviant.

The picture I am advocating is highly controversial, of course; but at least it is neat. It is organized around two main distinctions:

(i) The embedded sentence in a belief report contains either a singular term, or a quantified/descriptive phrase.

(ii) A quantified/descriptive phrase can be given either wide scope or narrow scope vis-a-vis the epistemic operator.

Thus there are three possibilities: what occurs in the embedded sentence can be a singular term, a quantified/descriptive phrase with narrow scope, or a quantified/descriptive phrase with wide scope. The belief report counts as *relational* if, and only if, the embedded sentence contains either a singular term or a quantified/descriptive phrase with wide scope. Note that there remains a difference between the two types of case. When using a singular term, the speaker himself makes a singular statement about the individual object the belief is about. When using a descriptive/quantified phrase with wide scope, the speaker ascribes a singular belief, but she does not herself express a singular belief, or make a singular statement, about the individual object the ascribed belief is about.

	<i>Expressed belief</i>	<i>Ascribed belief</i>	
<i>Genuine singular term</i>	singular	singular	
<i>Quantified or descriptive phrase taking wide scope</i>	general	singular	relational reading
<i>Quantified or descriptive phrase taking narrow scope</i>	general	general	notional reading

Table 1

At this point two main objections spring to mind:

- If the above theory was correct, it would be always be possible to infer from 'John believes that t is F' that there is an x John believes to be F. But what about statements like (10)?

(10) My three-years old son believes that Santa Claus will come tonight

Since Santa Claus does not exist, there is no individual to whom my son is related in the manner required for singular belief. Hence from (10) we cannot infer 'There is an x such that my son believes that x will come tonight'. That is a counter-example to the theory.

- I claim that belief sentences with singular terms are not ambiguous, in contrast to belief sentences with quantifiers. But they are: the name can be either endorsed by the speaker as his own way of referring to whatever the belief is about, or ascribed to the believer. That is the same old *de re/de dicto* ambiguity which we have observed in the case of belief sentences with quantifiers.

The second objection is especially important; it is the main obstacle on the road to accepting the view I have just sketched. In the next section, I will argue that it rests on a confusion. Belief sentences with singular terms are indeed ambiguous between a 'transparent' and an 'opaque' reading, but that ambiguity is *distinct from*, indeed orthogonal to, the relational/notional ambiguity we have been considering so far. When the two ambiguities are confused under a singular heading (the so-called '*de re/de dicto*' distinction), the situation becomes intractable and leads one to despair. Once the ambiguities are kept apart, however, the apparently intractable problems disappear.

As for the first objection, it can be rebutted as follows. The reason why we can't infer ' $(\exists x)$ (my son believes that x will come tonight)' from 'My son believes that Santa Claus will come tonight' is the same reason why we can't infer (12) from (11).

(11) Santa Claus lives in the sky

(12) $(\exists x)$ (x lives in the sky)

So the objection is not a specific objection to the view that genuine singular terms behave as such in belief reports; rather, it is an objection to the view that fictional names such as 'Santa Claus' *are* genuine singular terms, subject to ordinary logical principles. Since that problem is a general problem, it is not incumbent on the attitude theorist to solve it.

There is, however, an important difference between a fictional statement like (11) and a statement like 'My son believe that Santa Claus will come tonight' or 'In the story Santa Claus lives in the sky' ('metafictional' statements, as Currie [1990] aptly calls them). The author of a fictional statement does not really make assertions, but only pretends to do so. Thus in (11) she only pretends to say of a certain person that he lives in the sky.^{xii} Since that it so, the failure of existential generalisation is unproblematic. (12) cannot really be inferred, because (11) was not really asserted. (Within the pretense, however, the inference goes through: the speaker pretends to be committed to (12), by pretending to assert (11).) In contrast, it seems that metafictional statements are serious and evaluable as true or false (Lewis 1978). Hence it is not obvious that the failure of existential generalization has the same source in both cases.

Despite what I have just said, it can be maintained that the author of a metafictional statement such as 'In the story, Santa Claus lives in the sky' is also pretending: she pretends to assert of someone that the story says he lives in the sky. Similarly for (10): the speaker pretends to assert of a given individual that her son believes he will come tonight. In neither case does the speaker really make that assertion, as there is no individual the story (or the child's belief) is about. By

pretending to do so, however, the speaker communicates something true about the story or about the child's belief — something which could be communicated literally only by means of a lengthy and cumbersome paraphrase (Walton 1990: 396ff; Crimmins forthcoming; see also Forbes 1996 for discussion of related issues).

A lot more needs to be said to flesh out this proposal. One must detail the mechanism of 'semantic pretense' through which one can, in a more or less conventional manner, convey true things by pretending to say other things. One must also show how fictional statements like (11) can be distinguished from metafictional statements in which, intuitively at least, it seems that a genuine (and true) assertion is made. If pretense is involved in both cases, it is not quite the same sort of pretense; the theory owes an account of how the two kinds connect up with each other (Recanati forthcoming). I do not intend to go into those complex issues here, even though Quine turns out to have been a pioneer in this area too.^{xiii} It is sufficient to have pointed out that a research programme exists to solve precisely the sort of problem that (10) raises, in a way which is consistent with the theory I have expounded concerning the behaviour of singular terms in attitude contexts.^{xiv}

§2.4 *The ambiguity of the de re/de dicto distinction*

In §2.2 I glossed the relational/notional distinction in terms of the points of view involved. I said that the description (or the quantification) is "endorsed by the speaker" in relational readings, while it is "ascribed to the believer" in the notional reading. Now it seems that — contrary to what I claimed — exactly the same distinction can be made with respect to belief sentences containing singular terms instead of descriptions or quantifiers. Thus (13) can be understood in two ways.

(13) Ralph believes that Cicero denounced Catiline

On the transparent interpretation, Ralph is said to have a belief concerning the individual Cicero. Since Cicero is Tully, (13) can be rephrased as (14):

(14) Ralph believes that Tully denounced Catiline

The transparent reading of sentences like (13) is often rendered by appealing to the exported form, as in (15):

(15) Ralph believes of Cicero that he denounced Catiline

But there is another interpretation of (13) and (14), an interpretation in which they are not equivalent and cannot be rendered as (15). This is the 'opaque' interpretation. On that interpretation, Ralph is said by (13) to have a belief such that he would assent to 'Cicero denounced Catiline', but not necessarily to 'Tully denounced Catiline'. On the opaque interpretation, the use of the name 'Cicero' (rather than 'Tully') to refer to Cicero is *ascribed to the believer*. On the transparent reading, the choice of the name is up to the speaker and does not reflect the believer's usage; that is why replacement of 'Cicero' by 'Tully' in (13) on the transparent interpretation does not induce a change in the ascribed belief.

Quine and many philosophers and linguists after him have jumped to the conclusion that a single distinction applies to belief sentences whether they contain singular terms or descriptive/quantified phrases. They have equated the relational/notional distinction talked about in previous sections and the transparent/opaque distinction I have just introduced for belief sentences with singular terms. Both are viewed as instances of the so-called '*de re/de dicto*' distinction. The exported form (15) is the mark of the *de re*. Belief sentences on the *de dicto* (opaque, notional) reading resist exportation, because the epistemic operator takes wide scope — it governs the embedded sentence in its entirety. On the *de re* reading, the epistemic operator takes narrow scope and governs only the predicate: the subject expression, be

it quantificational or referential, is endorsed by the speaker without being ascribed to the believer. That is the confused doctrine whose untenability led Quine and others to despair of the original relational/notional distinction.

In fact, there is a clear difference between the two distinctions — the relational/notional distinction, and the transparent/opaque distinction. Consider the notional reading of a belief sentence. In such a case the believer is said to believe that there is an object *x* with such and such properties; that does not entail that there actually is an object *y* such that the believer believes that of *y*. Whatever quantification there is is strictly internal to the ascribed belief; it is not endorsed by the speaker. But *even on the opaque reading of a belief sentence in which a singular term occurs, reference is made to some particular individual* (Loar 1972). Thus the speaker who utters (13) on its opaque reading is committed to there being an individual *x*, such that Ralph's belief concerns *x* and is true iff ...*x*... To be sure, the belief which is ascribed to Ralph on the opaque reading of (13) is not merely the belief that that individual denounced Catiline; that would correspond to the transparent reading of (13). On the opaque reading, Ralph is ascribed the belief that: *Cicero* denounced Catiline. Cicero is thought of by Ralph not only as having denounced Catiline, but also *as Cicero*. Yet that feature of opacity is compatible with the relational character of the belief report, that is, with the fact that the speaker himself refers to Cicero as the object the ascribed belief is about. We can represent the opaque reading of (13) as follows:

(16) Ralph believes of Cicero, thought of as 'Cicero', that he denounced Catiline

The apposition 'thought of as *Cicero*' is sufficient to distinguish the opaque reading from the transparent reading. Both readings are relational: in both cases Ralph believes something of Cicero, and the speaker himself refers to Cicero as what Ralph's belief is about. In the opaque reading, however, the name has a dual role: it serves not only to refer to the object the ascribed belief is about, but also tells us something about how the

believer thinks of that object. As Brian Loar pointed out, this dual role is reminiscent of that of 'Giorgione' in Quine's famous example (Loar 1972: 51).

The non-equivalence of (13) and (14) on their opaque readings is clearly compatible with the relational character of those readings. In the same way in which (13), on its opaque reading, is rendered as (16), the opaque reading of (14) can be rendered as (17):

(17) Ralph believes of Cicero, thought of as 'Tully', that he denounced Catiline

The name 'Tully' in (14) refers to Cicero even on the opaque reading. The speaker is therefore committed to there being an individual, namely Cicero (= Tully), such that Ralph believes of that individual, thought of as 'Tully', that he denounced Catiline. There is no such existential implication when a belief report (with a descriptive or quantified phrase) is understood notionally.

As we can see, the contrast between cases in which something is ascribed to the believer and cases in which it is endorsed by the speaker is not drawn in quite the same way for the two distinctions. On the notional reading of a belief sentence with a descriptive/quantified phrase, the quantification is ascribed to the believer *without* being endorsed by the speaker; but the reference to the object of belief, and the existential commitment that goes with it, is *both* ascribed to the believer *and* endorsed by the speaker on the opaque reading of a singular belief sentence. The relational/notional distinction articulates a simple contrast between the point of view of the speaker and the point of view of the believer; while the transparent/opaque distinction articulates a quite different contrast, between the point of view of the sole speaker and the point of view of *both* the speaker and the believer. As far as the respective points of view of the speaker and the believer are concerned, opaque readings are thus essentially 'cumulative'.

Far from being identical to the relational/notional distinction, the transparent/opaque distinction turns out to be a distinction between two sorts of

relational reading. Hence there is no incompatibility between the claim that belief sentences with singular terms can only be understood relationally, and the observation that they have both a transparent and an opaque reading. Yet, precisely because belief reports with genuine singular terms cannot be interpreted notionally, but only relationally, it has seemed to many that a single distinction applies indifferently to all belief sentences: just as belief sentences with descriptive/quantified phrases can be interpreted relationally or notionally, belief sentences with singular terms can be interpreted transparently or opaquely. To dispell that illusion, one has only to notice that belief sentences with descriptive/quantified phrases are subject to *both* ambiguities. They can be interpreted notionally or relationally; and when relational, they can be interpreted transparently or opaquely. Loar gives the following example of a belief sentence with a quantified phrase which is naturally given a relational yet opaque interpretation:

(18) Ralph believes that a certain cabinet member is a spy

This does not mean that Ralph has a general belief to the effect that some cabinet member or other is a spy. As the phrase 'a certain' is meant to indicate, there is a particular cabinet member Ralph's belief is about. The belief report, therefore, is relational. However, Loar (1972: 54) points out that (18)

will often be taken to imply more than

(19) $(\exists y) (y \text{ is a cabinet member} \ \& \ B(\text{Ralph}, "x \text{ is a spy}", y))$

Ralph, we may suppose, believes it of the fellow under a certain description; that is,

(20) $(\exists y) (y \text{ is a cabinet member} \ \& \ B(\text{Ralph}, "x \text{ is a cabinet member and } x \text{ is a spy}", y))$

Loar's rendition of (18) as (20) nicely captures the cumulative aspect of opaque readings. Both the speaker and the believer view the person the belief is about as a cabinet member. As Loar pointed out (1972: 54), in a framework such as Quine's, in

which the two distinctions are conflated under a single heading, one cannot account for belief reports which, like (18), are both relational and opaque. 'Relational' entails 'transparent', for Quine and his followers. For that reason also, examples like (13) and (14), on their 'opaque' interpretation (corresponding to [16] and [17]), will have to be considered 'notional', while they are clearly relational. Given the extreme confusion that results, it is only natural that Quine eventually gave up the distinction as hopeless. It *is* hopeless, considered as a single distinction covering all the cases.

III. Opacity in belief sentences

§3.1 *'That'-clauses as complex demonstratives*

According to Brian Loar, the singular term in an opaque belief report has a dual role. It refers to the object the belief is about, but also determines an aspect of the ascribed belief concerning that object. The ascribed belief is conjunctive, and the first conjunct depends upon the identity of the singular term.

This theory can be understood in two ways. On one interpretation a singular term behaves deviantly in belief contexts. Instead of merely referring to some object, as singular terms normally do, it refers to an object *and* contributes a 'mode of presentation' to the content of the ascribed belief. That theory gives up semantic innocence, even if it does so in a less extreme manner than Frege's. It construes the singular term in a belief report as referential, but not as *purely* referential.^{xv}

There is another option, though. It consists in preserving semantic innocence and holding that the singular term in an opaque belief report is purely referential, in accordance with its normal semantic function. The opacity of the occurrence can then be explained by construing the context as *reflecting*, in analogy with the above analysis of the 'Giorgione' example (§1.2-3).

A context for a singular term is reflecting if and only if it contains a *dependent expression*, that is, an expression whose semantic value depends upon the identity of

the singular term occurring in that context. In the 'Giorgione' example, the dependent expression was the adverb 'so' in 'so-called', which we can construe either as demonstrative or as anaphoric. When we replace 'Giorgione' by a coreferential term, e.g. 'Barbarelli', the semantic value of the dependent expression changes. That accounts for the sentence's change in truth-value.

In the 'Giorgione' example, the dependent expression ('so-called') is part of the *frame* in which the singular term occurs ('... is so-called because of his size').^{xvi} The dependent expression is therefore disjoint (separable) from the singular term itself. But that need not be the case: for a context to be reflecting, it is not necessary that the dependent expression occur as part of the frame, in disjunction from the singular term itself. There are cases in which the singular term itself will be a constituent of the dependent expression. Let me give an example involving, not a singular but a general term.

Consider the demonstrative phrase 'that nag'. The semantic value of 'nag' is the same as that of 'horse'; the difference, as Frege would say, is one of 'colouring' rather than a properly semantic (truth-conditional) difference. Despite their semantic equivalence, 'nag' in 'that nag' cannot be replaced by 'horse', because the reference of a demonstrative phrase is linguistically underdetermined and crucially depends upon the referential intentions of the speaker, as revealed by the context. Now one aspect of the context which may be relevant to the determination of the speaker's referential intentions is the word which the speaker uses. If he uses a word such as 'nag', that provides some evidence that he does not intend to refer to his beloved and much respected horse Pablo, who happens to be otherwise salient in the context, but rather to the deprecated Pedro. If the word 'horse' was used, however, sheer salience would presumably promote Pablo to the status of referent. Substituting 'horse' for 'nag' can therefore change the likely interpretation, hence possibly the truth-value, of the sentence, by affecting the semantic value of the demonstrative phrase.

In general, whenever the semantic value of a phrase is linguistically underdetermined, and depends upon the intentions of the speaker, that phrase is a

reflecting context for its constituents. A 'part' of the global phrase cannot be replaced by a semantically equivalent expression without possibly affecting the semantic value of the whole, because any aspect of the context, including the actual words which are used, may be relevant to determining that semantic value.

Let us now go back to belief reports. If singular terms in belief sentences fulfill their ordinary function and are purely referential, substitutivity failures must be accounted for by appealing to the notion of a reflecting context. That means that we must find a dependent expression in the belief report — an expression whose semantic value depends upon the identity of the singular term.

One possible candidate is the 'that'-clause itself. A 'that'-clause is commonly taken to be a referring expression. Let us call what a 'that'-clause allegedly refers to a 'proposition', without going into the issue of what propositions exactly are. (I will construe them, heuristically, as $\$$ sentences in the sense of Kaplan 1986, since propositions in that sense seem to have been found palatable by Quine.) The reason why 'that'-clauses are generally considered as singular terms is that this enables us to account for inferences like the following:

John said that grass is green
 Everything John says is true
 Therefore, it is true that grass is green

If we rephrase 'It is true that grass is green' as 'That grass is green is true', as we are certainly entitled to do, the inference can easily be accounted for on the assumption that 'that grass is green' is a singular term. The pattern is:

a is F	(That grass is green is said by John)
Every F is G	(Everything said by John is true)
Therefore, a is G	(That grass is green is true)

Most philosophers consider the reference of a 'that'-clause as fixed by the following rule: a 'that'-clause refers to the proposition expressed by the embedded

sentence. In my book *Direct Reference* (Recanati 1993) I put forward an alternative proposal, in order to account for the well-known context-sensitivity of 'that'-clauses. I claimed that a 'that'-clause can, but need not, refer to the proposition expressed by the embedded sentence. It can also refer to a proposition obtained by contextually *enriching* the expressed proposition.

The relevant notion of contextual enrichment is that needed to account for examples like the following:

- (1) She took out her key and opened the door

In that example, analysed in Carston (1985), the fact that the door was opened with the key is not linguistically specified, yet it is certainly part of what we understand when we hear that sentence. It is an aspect of the meaning or content of the utterance which is provided through 'contextual enrichment'. John Perry calls that an 'unarticulated constituent' of what is said (Perry 1986); and he and Crimmins hold that modes of presentation of the objects of belief are unarticulated constituents of the proposition expressed by opaque belief reports (Crimmins and Perry 1989).^{xvii} I agree with the spirit, if not the details, of that analysis.

In my book I took a 'that'-clause to be a demonstrative phrase whose reference is constrained, but not determined, by the proposition which the embedded sentence expresses — much like the reference of the demonstrative phrase 'that horse' is constrained, but not determined, by the general term it contains. In other words, I took the reference of 'that'-clauses to be linguistically *underdetermined*. Underdetermination is to be distinguished from mere context-dependence. The reference of words like 'I' or 'today' is context-dependent, but it does not exhibit the relevant feature of underdetermination. In a given situation, the meaning of a pure indexical like 'I' or 'today' fully determines what the reference is. Not so with demonstratives. The reference of 'he' or 'that' is not determined by any rigid rule; it is determined by answering questions such as, Who or what can the speaker plausibly be taken to be

referring to, in that context? The same thing holds, I assumed, for 'that'-clauses. A 'that'-clause refers to a proposition which *resembles* the proposition expressed by the embedded sentence, but need not be identical with it; it can be an enrichment of it. What the reference of a given 'that'-clause actually is will depend upon the speaker's intentions as manifested in the context.

On that theory, when a belief report such as 'Ralph believes that Cicero denounced Catiline' is understood opaquely, the reference of the 'that'-clause 'that Cicero is a Roman orator' is distinct from what it is on the transparent interpretation. On the transparent interpretation the reference is, arguably, the 'singular proposition' (or valuated formula) which the embedded sentence expresses, viz. a sequence whose first member is the individual Cicero, and whose second member is the predicate 'Roman orator'. On the opaque interpretation, the reference is a 'quasi-singular' proposition, that is, the same thing except that the first member of the sequence is itself an ordered pair, consisting of the individual Cicero *and another predicate serving as 'mode of presentation'* (Recanati 1993). The quasi-singular proposition is an enrichment of the expressed singular proposition. The extra constituent provided by the context is the mode of presentation (the predicate) under which the reference of the singular term is assumed to be thought of by the believer.

The 'that'-clause thus turns out to be a dependent expression, whose semantic value is susceptible to change if a singular term occurring in the 'that'-clause is replaced by a coreferential term. For the reference of the 'that'-clause ultimately depends upon the speaker's communicative intentions as revealed by the context; and any aspect of the context, including the words which the speaker actually uses to report the ascriber's beliefs, may be relevant in figuring out the speaker's referential intention. In some contexts, the speaker's use of the name 'Cicero' will suggest that the believer thinks of Cicero as 'Cicero'. That is no more than a contextual suggestion, accountable perhaps in Gricean terms (McKay 1981; Salmon 1986); yet it may influence the assignment of a particular semantic value to the 'that'-clause, thereby affecting the truth-conditions of the belief report. That will be so whenever the belief report is understood opaquely: the

'that'-clause will then refer to a quasi-singular proposition involving not only the individual Cicero and the predicate 'Roman orator', but also a further predicate such as 'called *Cicero*'.

§3.2 *Semantic innocence*

Substitutivity problems have led many philosophers to give up semantic innocence in connection with attitude contexts. Both Frege and Quine thus appeal to the thesis of Semantic Deviance, according to which the extension of an expression is affected when it is embedded within a 'that'-clause. For Frege, the extension of a sentence systematically shifts in such circumstances. Once embedded, a sentence no longer denotes its truth-value but it comes to denote its truth-*condition*. Quine does not accept that view, but he sticks to the thesis of Semantic Deviance. 'That'-clauses, he says, are similar to quotation contexts: when we put a sentence in quotes, it no longer represents what it ordinarily represents. The sentence is mentioned rather than used.

In the case of quotation the thesis of Semantic Deviance is indeed very plausible. Do the first word of (2) and (3) below have e.g. the same extension? No, the first one refers to cats, the second one refers to the word 'cats'.

- (2) Cats are nice
- (3) 'Cats' is a four-letter word

But what about belief sentences and 'that'-clauses in general? Consider (4)

- (4) John believes that grass is green

Is it credible to say that the words 'Grass is green' do not represent what they normally represent? Does 'grass' in the embedded sentence refer to anything else than grass? Does it do anything else than refer to grass? As Davidson emphasized,

If we could recover our pre-Fregean semantic innocence, I think it would seem to us plainly incredible that the words 'The earth moves', uttered after the words 'Galileo said that', mean anything different, or refer to anything else, than is their wont when they come in different environments. (Davidson 1968: 144)

I fully agree with Davidson that we should at least *try* to "recover our pre-Fregean innocence", that is, to do without the thesis of Semantic Deviance. In an 'innocent' framework, the semantic value of an expression in the embedded part of a belief report is construed as its *normal* semantic value (whatever that may be).

I have shown how it is possible to preserve semantic innocence by construing the singular terms in belief contexts as purely referential (as they normally are), and accounting for failures of substitutivity in terms of reflecting context. Yet the theory elaborated in Recanati (1993), and summarized in §3.1, is not *thoroughly* innocent. It is innocent as far as singular terms are concerned, but when it comes to the complete embedded sentence, innocence is eventually abandoned.

Let us call the view that 'that'-clauses are complex singular terms the 'standard analysis'. On that view, the sentential complement *names* a proposition. But that is not what the complement sentence does when it is not embedded. Unembedded, the sentence expresses a proposition, it doesn't name one. Hence, by construing that-clauses as names, it seems that the standard account violates semantic innocence.

Faced with that objection, the usual strategy consists in drawing a distinction between the embedded sentence and the complete 'that'-clause. The embedded sentence, it is said, expresses a proposition, and it is that proposition which the 'that'-clause names. In this way innocence is allegedly saved: the sentence does the same thing — it expresses a certain proposition — whether it is embedded or not; it never names a proposition, since that is a job for the complete 'that'-clause.

I do not think this strategy works, however. First, the distinction between the embedded sentence and the complete 'that'-clause has no obvious equivalent when we turn to non-standard belief sentences like 'In John's mind, grass is green' or 'According

to John, grass is green'. There is no 'that'-clause in such examples — only the sentence 'Grass is green'. Second, even when the distinction makes syntactic sense, it is unclear that it enables us to preserve semantic innocence. I will show that by considering, once again, the case of quotation.

Faced with an instance of quotation such as (3), we have two options. We can say that the word 'cats' in this context does something different from what it normally does: it is used 'autonomously' (self-referentially). Or we can say that it is the complex expression consisting of the word 'cats' *and the quotes* which denotes the word 'cats'. If, by taking the second route, we refrain from ascribing the word 'cats' a deviant function in quotation contexts, we will be led to deny that the word 'cats' really occurs; rather, with Tarski and Quine, we will say that it occurs there only as a 'fragment' of the longer expression, much as 'cat' occurs in 'cattle'. From the semantic point of view, the relevant unit is indeed the complete quotation; the word 'cats' itself thus disappears from the picture. *In this way innocence is lost as surely as it is when we take the first option.* A truly innocent account is one that would *both* acknowledge the occurrence of the expression at issue in the special context under consideration *and* ascribe it, in that context, its normal semantic function. (Of course, there is no reason to expect an account of quotation to be semantically innocent in that sense.)

Similarly, we have two options with regard to attitude reports, in the standard framework. If we say that the complement sentence, once embedded, names the proposition which it would normally express, we give up semantic innocence: we accept that the embedded sentence does not do what it normally does. On the other hand, if, in order to protect innocence, we draw a sharp distinction between the embedded sentence (which expresses a proposition) and the 'that'-clause (which names it), *we run the risk of making the former disappear from the logical scene.* For the relevant semantic unit is the complete 'that'-clause. At the level of logical form the sentence 'John believes that S' has the form aRb — it consists of a two-place predicate and two singular terms. The embedded sentence plays a role only via the 'that'-clause in which it occurs. Which role? Arguably a *pre-semantic* role analogous to that of the

demonstration which accompanies a demonstrative. If that is right, then semantically the complexity of the 'that'-clause matters no more than the pragmatic complexity of a demonstrative-*cum*-demonstration or the 'pictorial' complexity of a quotation.^{xviii}

For that reason, I conclude that any theory which construes 'that'-clauses as singular terms is bound to give up semantic innocence at some point or other. On the view presented in §3.1, we protect the innocence of singular terms, but not that of the embedded sentence. The theory we end up with is thus not very different from Quine's. In Quine's framework, there is a sense in which the singular term 'Cicero' is purely referential in the sentence 'Cicero is bald' which we find embedded in the belief report 'John believes that Cicero is bald'. But the belief context into which the sentence is embedded is said by Quine to be 'opaque', like a quotation context. The situation is similar to that of:

(5) 'Cicero is bald' is held by John

The singular term 'Cicero' occurs purely referentially in the sentence 'Cicero is bald', but that sentence is quoted, hence insulated from the outer context by the opaque barrier of quotation. The pure referentiality of 'Cicero' thus becomes "strictly an internal affair" (Quine 1995: 356). At the outer level the insulated singular term no longer counts as purely referential. That is indeed Quine's definition of an opaque context: a context is referentially opaque "when, by putting a statement Φ into that context, we can cause a purely referential occurrence in Φ to be not purely referential in the whole context" (Quine 1953: 160).

A truly innocent account must give up the view that 'that'-clauses are singular terms. That is, the embedded sentence must be treated as, logically, a sentence; it must not be converted into a term. That means that 'Grass is green' has the same status in 'John believes that grass is green' which it has in 'The sky is blue and grass is green' (extensionality aside). The statement that grass is green is not mentioned, but simply

occurs as a part of a longer statement. (On the distinction between mentioning and compounding, see Quine 1962: 38.)

There are two theories on the market which satisfy this requirement of semantic innocence.^{xix} One is Davidson's paratactic theory (Davidson 1968); the other is Quine's sentential operator analysis, pursued and elaborated by Arthur Prior in a number of writings (Quine 1960; Prior 1963, 1971; Orenstein forthcoming). On this view, 'believes that' is an 'attitudinative' (Quine) or 'connecticate' (Prior) which forms a sentence from a singular term and a sentence. When its first argument place is filled by a singular term, it yields a sentence-forming operator, e.g. 'John believes that...'

The 'John believes that' operator can be viewed as a *world-shifting operator* (Recanati, forthcoming). It presents the sentence which follows it as true in John's belief world, rather than in the actual world. Whichever expression is responsible for such a world-shift is taken by Quine to constitute "an opaque interface between two ontologies, two worlds: that of the attitudinist, however benighted, and that of our responsible ascriber" (Quine 1995: 356). If Quine is right, then it does not matter whether we treat the embedded statement as mentioned, or as falling in the scope of a sentence-forming operator. In both cases the context in which the sentence is embedded is opaque. But is Quine right? If what I said concerning singular terms in belief contexts is correct, then, *pace* Quine, the ontology remains that of the ascriber all along, even though the 'world' which is described is that of the attitudinist: the objects the ascriber's belief is said to be about are picked out in the speaker's world, that is, in the actual world. If that view is tenable, there are two sorts of world-shift. One can use the singular terms with their normal references to describe counterfactual possibilities — worlds other than the actual world; let us call that an *innocuous* world-shift. Or one can 'shift ontologies' and use the terms with deviant references or at least without their normal references. In attitude contexts we have the first type of world-shift, but it is controversial that we have the more radical sort as well. The truly innocent theory I would like to see developed is one in which 'believes that' is an attitudinative, and the sentential operators built from it are innocuous world-shifters.

§3.3 *The attitudinatives as dependent expressions*

According to Hintikka (1962: 138-141), failures of substitutivity in belief contexts show that two co-referential singular terms, though they pick out the same individual in the actual world, may refer to different objects in the ascriber's belief world. That option is ruled out in the present framework; for we want the ontology to remain that of the ascriber all along: we want the singular terms to refer to the same objects, whether we are talking about the actual world or about the ascriber's belief world. That is the price to pay for semantic innocence. How, then, can we account for substitutivity failures?

Once again, we must appeal to the notion of a reflecting context. Consider the following inference:

- (6) John believes that Emile Ajar wrote *La Vie devant soi*
- (7) Emile Ajar = Romain Gary
- (8) John believes that Romain Gary wrote *La Vie devant soi*

Despite the identity stated in (7), we cannot infer (8) from (6). For it is possible that (6) is true and (8) false. That entails that (6) and (8) have different truth-conditions. Now, by virtue of (7) and the semantics of singular terms, the embedded sentence in (6) and (8) make the same (narrow) contribution to the truth-conditions of the global belief report. Hence it must be the interpretation of the prefix, 'John believes that', that is, *its* contribution to the truth-conditions of the global belief report, which changes from (6) to (8). If the prefix was given the same interpretation in (6) and (8), there could be no difference of truth-value between (6) and (8), since the embedded sentences express the same proposition. The prefix, therefore, must be a 'dependent expression' whose semantic value shifts as a result of the substitution.

To emphasize the similarity with Quine's 'Giorgione' example, and borrowing an idea from Graeme Forbes (1990), I suggest that we rephrase (6) and (8) respectively as

(6') John so-believes that Emile Ajar wrote *La Vie devant soi*

and

(8') John so-believes that Romain Gary wrote *La Vie devant soi*

In general, I suggest that whenever an attitude sentence, '*a* Ψs that *p*', is interpreted opaquely, we render it as '*a* so-Ψs that *p*', where 'so' is a demonstrative adverb referring to some manner of Ψ-ing instantiated in the context. Slightly more colloquially, we might use the phrase '*a* Ψs that *p* *thus*', or '*a* Ψs that *p* *in that manner*'. For example, '*a* says that *p*', opaquely understood, will be interpreted as tacitly referring to some manner of saying that *p*, as if the speaker had said: '*a* said that *p* *thus*'. Similarly for '*a* believes that *p*' and the other attitude verbs.

What is a manner of Ψ-ing? Consider the case of 'saying that'. Someone can say that I am ill by uttering the sentence 'He is ill' (while pointing to me) or by uttering 'Recanati is not well'. Those are two ways of saying that I am ill. Similarly, there are different ways or manners of believing that I am ill: by mentally entertaining the thought 'That guy is ill' or by entertaining the thought 'Recanati is not well'. As is well known, the distinction between a sentence and what it says extends to thought, and the corresponding distinction between *what* is believed and *how* that is believed provides a key to the puzzles of cognitive significance (Perry 1992).

Let us assume that the speaker utters

(9) John said that Recanati is not well

and that this is understood opaquely, as somehow reporting (some of) the words which John himself used. I analyse (9) as

- (9') John so-said that Recanati is not well
 = John said that Recanati is not well *thus*

where the demonstrative adverb, 'so' or 'thus', refers to some manner of saying that I am ill. Which manner of saying that? *The manner of saying which is instantiated by the speaker's utterance of the embedded sentence.*

In that framework the same prefix 'John believes that' makes different semantic contributions in (6) and (8), because the semantic value of the implicit demonstrative shifts when we substitute 'Romain Gary' for 'Emile Ajar'. The difference can be made explicit as follows:

(6'') John believes that Ajar/Gary wrote *La Vie devant soi* in that manner: 'Emile Ajar wrote *La Vie devant soi*'

(8'') John believes that Ajar/Gary wrote *La Vie devant soi* in that manner: 'Romain Gary wrote *La Vie devant soi*'

In (6'') and (8''), the adverbial 'in that manner' must of course be interpreted as modifying the main verb 'believes'. That verb itself must be given the 'transparent' interpretation: in (6'') and (8'') 'believes' is *not* equivalent to 'so-believes'.

As it stands the analysis is not wholly satisfactory, for not all aspects of the embedded sentence need to play a role in the imputation of a particular manner of believing to the ascriber. To refine the analysis, we can appeal to Nunberg's useful distinction between the *index* and the *referent* of a given occurrence of a demonstrative (Nunberg 1993). The index is what Kaplan (1989) calls the 'demonstratum' — that which is actually pointed to or attended at — but at least in cases of 'deferred ostension' that is distinct from the referent: the referent is the intended object, identifiable in relation to the index. Thus if, pointing to a car key, I say 'This is parked out back', the index (demonstratum) is the key, but the referent is the car. If we apply this distinction to our present case, we will say that the implicit demonstrative 'so' or 'thus'

demonstrates the speaker's current utterance of the embedded sentence (= index), and thereby *refers to* a certain manner of Ψ -ing, namely, that manner of Ψ -ing which would be instantiated if one Ψ -ed by uttering/entertaining that sentence.^{xx} On this analysis not all aspects of the demonstrated utterance need to be relevant to the determination of the manner of Ψ -ing which the speaker ascribes to the believer.

We can achieve the same result without appealing to Nunberg's distinction, however. Instead of analysing '*a so-believes that p*' as '*a believes that p in that manner*', we can analyse it, more perspicuously perhaps, as: '*a believes that p like that*', where the demonstrative '*that*' refers to the utterance of the embedded sentence. The manner of Ψ -ing denoted by the whole adverbial phrase '*like that*' will then depend upon the dimensions of similarity which are contextually relevant.

Whichever method we choose, the prefix turns out to be context-sensitive in two distinct ways, on the opaque interpretation. First, its semantic value depends upon the embedded sentence which follows it; for the demonstratum *d* (the index, or the referent of the constituent demonstrative '*that*') automatically changes when we substitute an expression for another in the embedded sentence. That, in itself, is sufficient to account for failures of substitutivity in attitude contexts. Second, the manner of Ψ -ing *m* which the demonstrated utterance is taken to instantiate will itself depend upon the aspects of the demonstrated utterance which are considered relevant. Even if we fix the demonstrated utterance, it will still be possible, by changing the context, to change the manner of Ψ -ing ascribed to the Ψ -er, thereby affecting the semantic value of the prefix.

There is, of course, an even more basic dimension of contextual variation: the belief report can be understood as transparent or opaque in the first place. The opaque reading I take to be a contextual enrichment of the transparent reading. Much as '*She opened the door*' in (1) is contextually enriched into '*She opened the door with the key*', '*John believes that p*' is enriched into '*John believes that p in such and such manner*'. The transparent/opaque ambiguity for belief reports is therefore an ambiguity between the minimal reading and a contextually enriched reading of the sentence.^{xxi} Here as

elsewhere, the enriched reading entails the minimal reading. (That is a general property of enrichment. See Recanati 1993.)^{xxii}

Table 2 summarizes the three dimensions of contextual variation we have discerned in belief sentences.

Minimal vs contextually enriched interpretation

Variation of the demonstratum

Variation of the manner of Ψ-ing taken to be instantiated by the demonstratum

transparent (minimal) reading:
a believes that *p*

believes that *p* like *d*₁

believes that *p* in manner *m*₁
 believes that *p* in manner *m*₂
 believes that *p* in manner *m*₃
 believes that *p* in manner *m*₄

....

believes that *p* like *d*₂

believes that *p* in manner *m'*₁
 believes that *p* in manner *m'*₂
 believes that *p* in manner *m'*₃
 believes that *p* in manner *m'*₄

...

opaque (enriched) reading:
a so-believes that *p*
 = *a* believes that *p* like that

believes that *p* like *d*₃

believes that *p* in manner *m''*₁
 believes that *p* in manner *m''*₂
 believes that *p* in manner *m''*₃
 believes that *p* in manner *m''*₄

...

...

...

...

...

...

...

Table 2

I conclude, first, that Quine was quite right to stress the extreme context-sensitivity of attitude reports; second, that the content of the embedded sentence need not be considered as affected by the contextual variation. All the shifts in interpretation talked about in this section can be construed as changes in the semantic value of the prefix 'a Ψ s that'. It is the prefix which can be interpreted minimally or in an enriched, opaque manner ('a so- Ψ s that'), depending on the context;^{xxiii} and it is the semantic value of the prefix which, on the opaque interpretation, varies according to the two further sorts of contextual change I have described.

§3.4 *Opacity, substitution, and quantification*

The prefix 'John believes that' is a dependent expression only on the opaque reading, i.e., when it is interpreted as 'John so-believes that'. On the transparent reading it is not a dependent expression. Since the occurrences of singular terms in the embedded sentences are uniformly treated as purely referential, in accordance with their normal semantic function, they come out *transparent*, by the definitions given in part I, whenever the prefix itself is given the transparent reading: for (i) they are purely referential, and (ii) the context in which they occur is not reflecting (since the prefix is not a dependent expression, on the transparent reading). The truth-value of a transparent belief report therefore depends only upon the reference of the term, not on its identity. That strongly suggests that singular terms in transparent belief contexts should be substitutable, that is, freely replaceable by coreferential singular terms. Yet, I shall argue, they are not.

The reason why occurrences of singular terms in transparent belief reports are not substitutable, even though they are transparent, is very simple. Since (i) the prefix can be given an opaque (enriched) as well as a transparent (minimal) interpretation, depending on the context; and (ii) substituting an expression for another in the embedded sentence changes the context in which the prefix is tokened; it follows that the substitution can shift the interpretation of the prefix from transparent to opaque, by

making it more likely that the speaker, using *those* words, intends to capture the believer's own way of thinking of the matter. In other words, belief contexts are *unstable* (§1.5). Only if we somehow fix (stabilize) the transparent interpretation of the prefix will substitution of coreferentials be a legitimate move.

Our findings so far can be summarized as follows:

- An occurrence of a singular term in the embedded portion of a belief sentence is purely referential, but not necessarily transparent: it is transparent only if the belief sentence is given a minimal interpretation ('transparent' reading). When the belief sentence is given an enriched interpretation ('opaque' reading), the occurrence of the singular term is not transparent, because the context in which it occurs is reflecting.
- Whether transparent or not, an occurrence of a singular term in the embedded portion of a belief sentence is not substitutable. It is non-substitutable either because the context *is* reflecting (opaque reading) or because the substitution can *make it* reflecting (transparent reading).

A last feature of singular belief reports must now be considered. As we saw in the first part of this paper, Quine tends to equate pure referentiality, transparency and substitutability. There is a fourth, no less important property on Quine's list: existential generalizability. When a singular term is purely referential (transparent, substitutable), Quine says, existential generalization is possible. When substitutivity fails because of opacity, existential generalization likewise fails. Thus we cannot go from

- (10) Giorgione was so-called because of his size
to
(11) $(\exists x)$ (x was so-called because of his size)

Contrary to Quine, who holds that transparency entails substitutability, I emphasized that even transparent occurrences of singular terms in (transparent) belief reports are not substitutable — unless of course we stabilize the context by fixing the interpretation of all the other expressions in the sentence while we make the substitution. A first question that arises, therefore, is this: Is a transparent occurrence of a singular term in a belief context open to existential generalization? If the answer is, as I claim, 'yes', then, *pace* Quine, substitution and existential generalization do not go hand in hand. I will go much further than that: I will argue that even opaque occurrences of singular terms in belief contexts are open to existential generalization. On the picture I am advocating (Table 3), substitution is *never* possible (even if the occurrence of the singular term at issue is transparent); while existential generalization is *always* possible (even if the occurrence of the singular term at issue is opaque).

<i>Occurrence of in singular term</i>	<i>'opaque' belief report</i>	<i>'transparent' belief report</i>
<i>purely referential?</i>	yes	yes
<i>transparent?</i>	no	yes
<i>substitutable?</i>	no	no
<i>open to existential generalization?</i>	yes	yes

Table 3

The unstability of the context accounts for the (surprising) failure of substitutivity in transparent belief reports. Substitutivity fails because the substitution can, by changing the context, shift the interpretation of the prefix from transparent to opaque, thereby affecting the truth-conditions of the belief report. It is also the unstability of the context which accounts for the (no less surprising) possibility of existential generalisation in opaque belief reports.

Normally, opacity blocks existential generalization. For the truth-value of a sentence containing an opaque occurrence of a singular term depends upon the identity of the term, not merely on its reference. When that term is eliminated through existential generalization, the statement is left incomplete and unevaluable: a reflecting context with nothing to reflect. Thus, Quine observes, (11) "is clearly meaningless, there being no longer any suitable antecedent for 'so-called' " (Quine 1961: 145). There is, however, a crucial difference between an opaque belief sentence and a sentence like (10) — a difference which accounts for the success of existential generalization in opaque belief sentences.

(10) is a reflecting context for the singular term 'Giorgione', and it is so in a stable manner: the context *remains* reflecting under operations such as substitution of coreferentials or existential generalization. But a belief sentence is a reflecting context for the singular terms occurring in the embedded sentence only when it is given the opaque ('so-believes') interpretation; and that interpretation is a highly context-sensitive hence *unstable* feature of the sentence. As I have repeatedly stressed, replacing a transparent occurrence of a singular term by an occurrence of a coreferential singular term may change the truth-value of the report by shifting the prefix from the transparent to the opaque interpretation. In the other direction, replacing an opaque occurrence of a singular term, that is, an occurrence of a singular term in the embedded portion of an opaquely interpreted belief report, by a variable, automatically shifts the interpretation of the prefix from the opaque ('so-believes') to the transparent interpretation; for it is only on the transparent reading that the quantified statement makes sense. If the context remained reflecting, the statement would become meaningless once the singular term is eliminated. By virtue of this compensatory mechanism, we can go from 'Tom believes that Cicero denounced Catiline', even on the opaque interpretation, to 'Someone is such that Tom believes *he* denounced Catiline'. The opacity of the original sentence is pragmatically filtered out in the very process of existential generalization.

At this point one might argue that, surely, the inference is illegitimate. We can go by existential generalization from 'Fa' to ' $(\exists x) (Fx)$ ', but not from 'Fa' to ' $(\exists x) (Gx)$ '.

But in the type of inference I have just described, an expression (viz. the attitudinative) is interpreted differently in the premiss (the opaque belief sentence we start with) and the conclusion (the quantified statement). Logically, therefore, the inference does not take us from 'Fa' to ' $(\exists x) (Fx)$ ', but from 'Fa' to ' $(\exists x) (Gx)$ '. That is an instance of the fallacy of equivocation mentioned in §1.3.

But I think existential generalization from opaque belief reports with singular terms *is* a valid move. Since the opaque reading is an enrichment of the transparent reading, it entails the transparent reading. It is therefore legitimate to go from the opaque belief report 'Tom believes_o that Cicero denounced Catiline' to the meaningful quantified statement 'There is someone of whom Tom believed_t he denounced Catiline'; for the latter is entailed by the transparent belief report 'Tom believes_t that Cicero denounced Catiline', and that transparent belief report itself is entailed by the opaque belief report.

Appendix

Even though I consider 'Giorgione' as purely referential in Quine's famous example, I accept Quine's point that there are intermediate cases between pure autonymy and pure referentiality. A good example is

- (1) A 'robin' is a thrush in American English, but not in British English.

Though it is quoted, the word 'robin' here keeps its normal semantic value: it denotes a type of bird. It is a type of bird, not a word, which is said to be a thrush. But (1) also says something about the word 'robin'. For it is the word, not the bird, whose properties vary from one dialect of English to the next. As Austin pointed out, this mixture of mention and use is typical of semantic discourse:

Although we may sensibly ask 'Do we *ride* the word "elephant" or the animal?' and equally sensibly 'Do we *write* the word or the animal?' it is nonsense to ask 'Do we *define* the word or the animal?' For defining an elephant (supposing we ever do this) is a compendious description of an operation involving both word and animal (do we focus the image or the battleship?). (Austin 1971: 124)

'Echoes' provide another example of mixed use. Often one uses a word while at the same time implicitly ascribing that use to some other person (or group of persons) whose usage one is blatantly echoing or mimicking. Thus one might say:

(2) That boy is really 'smart'

In such examples one is quoting, but at the same time using the words with their normal semantic values.

In (2) the fact that a word is quoted while being used does not affect the truth-conditions of the utterance. But sometimes it does. Thus I can refer to some object, A, using the name of another object, B, in quotes, providing the person I am mimicking uses the name for B as a name for A. I may well say

(3) 'Quine' has not finished writing his paper

and refer, by the name 'Quine' in quotes, not to Quine but to that person whom our friend James mistakenly identified as Quine the other day. Any word can, by being quoted in this echoic manner, be ascribed a semantic value which is not its normal semantic value, but rather what some other person takes to be its semantic value.^{xxiv}

Many instances of mixed use lend themselves to a paraphrase where the expression in quotes is replaced by a descriptive phrase in which that expression occurs autonomously. Thus (1) could be rephrased as

- (1*) The bird called 'robin' is a thrush in American English,
but not in British English^{xxv}

The description 'the bird called 'robin'' describes a bird, but it does so by mentioning a word. We might, similarly, paraphrase the 'echoic' examples above using such metalinguistic descriptions.^{xxvi}

Footnotes

* My thanks to Brian Loar for helpful comments on this paper.

ⁱ I am using 'Wychnevetsky' as an arbitrary name for the word 'cat'.

ⁱⁱ To the extent that he makes that claim, Quine does not take it at face value. He does not say that the occurrence is accidental *simpliciter*, but that it counts as accidental *from a logical point of view*. For a gloss, see the end of this section.

ⁱⁱⁱ Indeed the expression in question may differ from the enclosed word in gender or number. Thus in:

'Cats' is a four-letter word

the first word of the sentence is singular, while the mentioned word ('cats') is plural.

^{iv} The three examples I have just given contrast with the previous examples (autonymy, metonymy, etc.) in the following respect: In autonymy or metonymy there is a distinction between the normal semantic value and the semantic values which deviate from the norm. In the type/token or process/product ambiguity, arguably there is no such asymmetry. (See Nunberg 1979 for an investigation of those issues.) Despite this difference, all the examples belong to the family of 'systematic ambiguities', whose study has proved important and fruitful in contemporary semantics (see e.g. Pustejovsky 1995).

^v The autonymous word need not be taken as self-referential, of course. Instead of having the word refer to itself, we can insist that it is the complex expression (word plus quotation marks) or the pair of quotation marks (construed as a demonstrative, in the manner of both Prior [1971: 60-61] and Davidson [1979]) which refers to the quoted word.

^{vi} In Quine's framework, the first type of case (the case in which it is the singular term itself which is used non-purely referentially) can be reduced to the second type of case (the case in which it is the context or the position that generates opacity). For Quine takes an autonomous word to be a word that occurs in a special linguistic context, viz. 'within quotation marks'. Quine can thus get rid of non-purely referential occurrences or uses altogether and handle opacity entirely in terms of positions. Opaque or non-purely referential positions are linguistic contexts (e.g. quotation contexts or reflecting contexts) such that a singular term in that context is not subject to the Principle of Substitutivity. In this way a uniform treatment is provided for

'...' is a three-letter word

and

... is so-called because of his size

Even so, there remains a big difference between the two types of case. In the first context the singular term is not used referentially (it is deviant), while it is used with its normal referential function in the second context.

^{vii} As Kit Fine has shown, we can make any non-referential occurrence of a singular term similarly substitutable merely by forming the disjunction of the sentence where it occurs with ' $2 + 2 = 4$ ' (Fine 1989: 218).

^{viii} See Forbes 1990; Crimmins 1992; Recanati 1993. The first after Quine himself to have drawn attention to the analogy between belief sentences and 'Giorgione' sentences was Brian Loar in his 1972 article.

^{ix} To say that they have a non-deviant referential use is compatible with saying that their referential use is not their normal semantic function. If their referential use was their normal semantic function, their non-referential uses would themselves be deviant. In Recanati (1993) I claimed that attributive and referential uses of descriptions are *both* non-deviant. (In contrast, only purely referential uses of singular terms are non-deviant.)

^x "If we say that there is someone of whom Othello believes that she is unfaithful, while we do not thereby put *ourselves* into any relation with anyone except Othello, we do thereby say that there is someone with whom *he* stands in the relation of believing her unfaithful" (Prior 1971: 135).

^{xi} In 1956 Quine said that exportation — the step from '*a* believes that *t* is *F*' to '*a* believes of *t* that it is *F*' — "should doubtless be viewed in general as implicative" (Quine 1956: 190). Afterwards he was moved by the Sleigh/Kaplan example of the shortest spy (Sleigh 1968; Kaplan 1968): if exportation is valid, then we can go from 'John believes that the shortest spy is a spy' to ' $(\exists x)$ (John believes *x* is a spy)', via 'John believes of the shortest spy that he is a spy'; but if that is accepted, an obviously notional belief report is treated as if it were relational. As Quine concludes, "we must find against exportation" (Quine 1977: 9). Indeed, insofar as exportation opens the way for existential quantification, it is clear that exportability must be restricted to those cases in which the relational reading is intuitively appropriate. It cannot be treated as generally permissible.

Still, I think Quine was right in the first place: exportation *is* generally valid, provided *t* is a genuine singular term. In the Sleigh/Kaplan example, it isn't. Of course, exportation also works when *t*, though not a singular term, is given scope over the epistemic operator. But that is not the case in Sleigh/Kaplan example either. In the Sleigh/Kaplan example Ralph is said to believe 'The shortest spy is a spy'; he is not said to believe, of some particular individual known to him (and described by the speaker as 'the shortest spy'), that he is a spy.

^{xii} (11) can also be interpreted as short for 'In the story, Santa Claus lives in the sky' (Lewis 1978). On that interpretation (11) is a metafictional statement, like (10).

^{xiii} See Quine's seminal and oft-quoted remarks on the role of empathy and pretense in belief ascriptions (Quine 1960: 219).

^{xiv} The theory in question should not be considered as dependent upon the success of the Walton-Crimmins research programme, however. Should the latter fail, it would

still be possible to argue that metafictional uses (e.g. the use of 'Santa Claus' in (10)) are somehow deviant, like the non-purely referential uses mentioned in the appendix.

^{xv} To flesh out this interpretation, one might construe the occurrence of a singular term in an opaque belief context as 'echoic', along the lines of

Ralph believes of 'Cicero' that he denounced Catiline

where the quotes around the singular term 'Cicero' indicate that the occurrence is partly autonomous. On echoic uses, see the Appendix below.

^{xvi} Brian Loar also uses the word 'frame' in his 1972 article, but in a different sense.

^{xvii} That is, in effect, the 'hidden indexical theory' of belief reports, the first formulations of which can be found in Linsky (1967: 113) and Schiffer (1977: 32-33).

^{xviii} "From the standpoint of logical analysis each whole quotation must be regarded as a single word or sign, whose parts count no more than serifs or syllables. A quotation is not a *description*, but a *hieroglyph*; it designates its object not by describing it in terms of other objects, but by picturing it. The meaning of the whole does not depend upon the meanings of the constituent words." (Quine 1951: 26)

^{xix} Three if we count the theory put forward in Panaccio 1996. According to Panaccio, 'John believes that S' can be analysed as 'John believes something which is true *tantumsi* S', where '*tantumsi*' is a special, indexical connective.

^{xx} This counterfactual circumlocution is necessary because the speaker herself need not be Ψ -ing in uttering the embedded sentence.

^{xxi} Some will hasten to conclude that only the transparent reading is relevant to semantics, since semantics is supposed to deal with literal meanings and 'minimal' interpretations, without considering 'speaker's meaning'. (See e.g. Salmon 1986.) I disagree. If we want our theory to keep in touch with our intuitions, we cannot disregard those aspects of meaning resulting from enrichment. Moreover, when we embed an utterance, we see that the aspects of meaning which result from contextual enrichment often become part of the minimal interpretation of the complex utterance,

thereby making the intended segregation of the semantic from the pragmatic untenable. For more on those matters, see Recanati 1993, part II.

^{xxii} At this point a suggestion must be considered: can we not reduce the third form of context-sensitivity to the second one? Remember that in the opaque reading, there are two dimensions of variation. (i) Any change in the wording of the embedded sentence can, by changing the demonstratum, affect the semantic value of the sentential operator; (ii) even if the demonstratum is fixed, the manner of Ψ -ing which is tacitly referred to can vary depending on the aspects of the index that are considered relevant in the context at hand. Now we could consider the transparent reading as the special case in which the relevant aspect of the demonstrated utterance is nothing other than its semantic content — its truth-conditions — to the exclusion of any other properties of the embedded sentence.

I have two objections to this move:

(i) Instead of considering the opaque reading as an enrichment of the transparent reading, it construes the transparent reading as a limiting case of the opaque. But in analysing the opaque reading in terms of 'so-believing' or 'believing thus', I implicitly accepted the primacy of the transparent sense. If we use subscripts to distinguish the transparent sense from the opaque, then, on my analysis, to believe_o that p is to believe_t that p in such and such a manner. That is hardly consistent with analysing believe_t that p as a variety of believe_o that p.

(ii) If we construe 'that'-clauses as singular terms, it is indeed tempting to view transparency as a limiting case of opacity. Thus we can say that a 'that'-clause always refers to some enrichment of the proposition expressed, where that proposition itself counts as the minimal or 'zero' enrichment. But such a move is much less tempting when, giving up the view that 'that'-clauses are singular terms, we construe the embedded sentence as a *bona fide* sentence. On that construal the semantic content of the embedded sentence (the proposition it expresses) is its contribution to the content of the global attitude report: the embedded sentence expresses a proposition, and the

sentential operator maps that proposition on a truth-value (truth, whenever the proposition in question holds in the ascriber's belief world). The semantic content of the embedded sentence is thus given prior to and independent of whatever tacit reference to some manner of Ψ -ing the sentential operator may additionally convey.

^{xxiii} The 'context' in which the prefix is tokened includes the words it is prefixed to. See §3.4 for some consequences of that type of context-sensitivity.

^{xxiv} In Recanati 1987: 63 I offered an example involving a definite description:

Hey, 'your sister' is coming over

In that example the person who is coming over is not actually the addressee's sister, but is thought to be the addressee's sister by some person whom the speaker is ironically mimicking.

^{xxv} Note that the description has to be construed as an 'incomplete' definite description, that is, as a definite description whose domain is contextually variable, on the pattern of: 'The window is open in the kitchen, but not in the living room.' (See Recanati 1996 for a treatment of incomplete definite descriptions in terms of contextually variable domains.)

^{xxvi} There is a clear difference, in all cases, between the original sentence and the explicit paraphrase; a difference which should prevent one from saying that one is 'elliptical for' the other. See Walton 1990: 222-224.

References

- Austin, J.L. (1971), *Philosophical Papers*, 2nd ed. Oxford: Clarendon Press.
- Barwise, J. and J. Perry (1981), Semantic Innocence and Uncompromising Situations. *Midwest Studies in Philosophy* 6:387-403.
- Carston, R. (1988), Implicature, Explicature, and Truth-Theoretic Semantics. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*, Cambridge: Cambridge University Press, p.155-81.
- Chastain, C. (1975), Reference and Context. In K. Gunderson (ed.), *Language, Mind, and Knowledge*, Minneapolis: University of Minnesota Press, p. 194-269.
- Crimmins, M. (1992), *Talk about Belief*. Cambridge, Mass: MIT Press/ Bradford Books.
- Crimmins, M. (forthcoming), Hesperus and Phosphorus.
- Crimmins, M. and J. Perry (1989), The Prince and the Phone Booth. *Journal of Philosophy* 86: 685-711.
- Currie, G. (1990), *The Nature of Fiction*. Cambridge: Cambridge University Press.
- Davidson, D. (1968), On Saying That. *Synthese* 19: 130-146.
- Davidson, D. (1979), Quotation. *Theory and Decision* 11: 27-40.
- Dever, J. (forthcoming), Strangers on a Train: Engineering a Kripke-Frege Reunion.
- Donnellan, K. (1966), Reference and Definite Descriptions. *Philosophical Review* 75:281-304.
- Evans, G. (1982), *The Varieties of Reference*, ed. by J. McDowell. Oxford: Clarendon Press.
- Fine, K. (1989), The Problem of De Re Modality. In J. Almog, H. Wettstein and J. Perry (eds.), *Themes from Kaplan*, New York: Oxford University Press, p. 197-272.
- Forbes, G. (1990), The Indispensability of Sinn. *Philosophical Review* 99:535-563.
- Forbes, G. (1996), Substitutivity and the Coherence of Quantifying In. *Philosophical Review* 105: 337-372.

- Geach, P.T. (1972), *Logic Matters*. Oxford: Basil Blackwell.
- Hintikka, J. (1962), *Knowledge and Belief*. Ithaca: Cornell University Press.
- Hornsby, J. (1977), Singular Terms in Contexts of Propositional Attitude. *Mind* 86: 31-48.
- Kaplan, D. (1968), Quantifying In. *Synthese* 19: 178-214.
- Kaplan, D. (1986), Opacity. In L. Hahn and P. A. Schilpp (eds.) *The Philosophy of W. V. Quine*, La Salle, Illinois: Open Court, p. 229-89.
- Kaplan, David (1989), Demonstratives. In J. Almog, H. Wettstein and J. Perry (eds.) *Themes from Kaplan*, New York: Oxford University Press, p. 481-563.
- Kripke, S. (1977), Speaker's Reference and Semantic Reference. *Midwest Studies in Philosophy* 2:255-276.
- Lewis, D. (1978), Truth in Fiction. *American Philosophical Quarterly* 15: 37-46.
- Linsky, L. (1967), *Referring*. London: Routledge & Kegan Paul.
- Loar, B. (1972), Reference and Propositional Attitudes. *Philosophical Review* 81:43-62.
- McKay, T. (1981), On Proper Names in Belief Ascriptions. *Philosophical Studies* 39:287-303.
- Neale, S. (1990), *Descriptions*. Cambridge, Mass: MIT Press/Bradford Books.
- Nunberg, G. (1979), The Non-Uniqueness of Semantic Solutions: Polysemy. *Linguistics and Philosophy* 3:143-84.
- Nunberg, G. (1993), Indexicality and Deixis. *Linguistics and Philosophy* 16: 1-43.
- Orenstein, A. (forthcoming), Propositional Attitudes without Propositions. In P. Kotatko and A. Grayling (eds.), *Meaning*, Oxford: Clarendon Press.
- Panaccio, C. (1996), Belief-Sentences: Outline of a Nominalistic Approach. In M. Marion and R.S. Cohen (eds.), *Québec Studies in the Philosophy of Science II*, Dordrecht: Kluwer, p.265-277.
- Perry, J. (1986), Thought without Representation. *Proceedings of the Aristotelian Society, Supplementary Volume* 60: 137-51.
- Perry, J. (1992), *The Problem of the Essential Indexical and Other Essays*. New York Oxford University Press.

- Prior, A. (1963), *Oratio Obliqua*. *Proceedings of the Aristotelian Society*, *Supplementary Volume 37*: 115-126.
- Prior, A. (1971), *Objects of Thought*, ed. by P. Geach and A. Kenny. Oxford: Clarendon Press.
- Pustejovsky, J. (1995), *The Generative Lexicon*. Cambridge, Mass: MIT Press.
- Quine, W.V.O. (1951), *Mathematical Logic*, 2nd ed. Cambridge, Mass: Harvard University Press.
- Quine, W.V.O. (1953), Three Grades of Modal Involvement. Reprinted in Quine 1976, p. 158-176.
- Quine, W.V.O. (1956), Quantifiers and Propositional Attitudes. Reprinted in Quine 1976, p. 185-196.
- Quine, W.V.O. (1960), *Word and Object*. Cambridge, Mass: MIT Press.
- Quine, W.V.O. (1961), *From a Logical Point of View*, 2nd ed. Cambridge, Mass: Harvard University Press.
- Quine, W.V.O. (1962), *Methods of Logic*, 2nd ed. London: Routledge & Kegan Paul.
- Quine, W.V.O. (1974), *The Roots of Reference*. La Salle, Illinois: Open Court.
- Quine, W.V.O. (1976), *The Ways of Paradox and Other Essays*, 2nd ed. Cambridge, Mass: Harvard University Press.
- Quine, W.V.O. (1977), Intensions Revisited. *Midwest Studies in Philosophy 2*: 5-11.
- Quine, W.V.O. (1980), The Variable and Its Place in Reference. In Z. van Straaten (ed), *Philosophical Subjects: Essays presented to P.F. Strawson*, Oxford: Clarendon Press, p. 164-173.
- Quine, W.V.O. (1995), Reactions. In P. Leonardi and M. Santambrogio (eds.), *On Quine*, Cambridge: Cambridge University Press, p. 347-361.
- Recanati, F. (1987), Contextual Dependence and Definite Descriptions. *Proceedings of the Aristotelian Society 87*:57-73.
- Recanati, F. (1993), *Direct Reference: From Language to Thought*. Oxford: Basil Blackwell.
- Recanati, F. (1996), Domains of Discourse. *Linguistics and Philosophy 19*: 445-475.

Recanati, F. (forthcoming), The Iconicity of Metarepresentations. In D. Sperber (ed.), *Metarepresentations*, New York: Oxford University Press.

Russell, B. (1905), On Denoting. *Mind* 14: 479-493.

Salmon, N. (1986), *Frege's Puzzle*. Cambridge, Mass: MIT Press/Bradford Books.

Schiffer, S. (1977), Naming and Knowing. *Midwest Studies in Philosophy* 2:28-41.

Sleigh, R.C. (1968), On a Proposed System of Epistemic Logic. *Noûs* 2: 391-398.

Walton, K.L. (1990), *Mimesis as Make-Believe*. Cambridge, Mass: Harvard University Press.

CREA, CNRS/Ecole Polytechnique
1 rue Descartes, 75005 Paris, France