

Modes of presentation in attitude reports

François Recanati

Forthcoming in A. Sullivan (ed.)
*Sensations, Thoughts, Language :
Essays in honor of Brian Loar*
London : Routledge

1

In ‘Quantifiers and Propositional Attitudes’ (1956), Quine emphasizes the importance of a distinction between two types of attitude reports : *de re* and *de dicto*, or, as Quine himself puts it, ‘relational’ and ‘notional’. The distinction is based upon a more fundamental distinction at the level of the attitudes themselves.

Some attitudes are relational : they involve a relation between the attitude holder and some other individual. For example, the attitude holder may believe, of his friend John, that he is always late. This is a (triadic) relation between the attitude holder, his friend John, and what the former believes of the latter (viz., that he is always late). It is natural to report that sort of belief (belief in the relational sense, as Quine puts it) by using the ‘believes of’ locution, as I have just done. Other attitudes are not relational in that sense : they do not involve a relation between the attitude holder and some other individual, but only (as far as relations go) a relation between the attitude holder and the propositional content of the attitude. For example, the attitude holder may believe that there are spies. This is a (dyadic) relation between the attitude holder and the proposition that there are spies.

The belief that there are spies is the belief that the property of being a spy is instantiated – the belief that someone or other is a spy. But if I say, ‘John believes that someone is a spy’, that is ambiguous. That can be read either as saying that, according to John, there are spies (nonrelational reading : John believes that the concept of spy is instantiated), or as saying that there is someone in particular *of whom* John believes that he is a spy (relational reading). Quine traces the ambiguity to the belief predicate, which can be dyadic or triadic. On the relational reading of ‘John believes that someone is a spy’, ‘believes’ is triadic and the second argument position (the person of whom John believes that he is a spy) is existentially quantified :

$$\exists x B^3 (\text{John}, x, \underline{z \text{ is a spy}})$$

On the nonrelational reading the belief is about no one in particular : it is a (dyadic) relation between John and a general proposition, to the effect that someone is a spy.

$$B^2 (\text{John}, \underline{\exists x x \text{ is a spy}})$$

The same sort of ambiguity arises if, instead of a quantifier such as ‘someone’, we use an indefinite description. Thus ‘John wants to marry a woman richer than himself’ can be understood in two ways. Understood one way the utterance entails that there is a particular

woman richer than himself whom John wants to marry. Understood the other way it does not : John's desire involves the *notion* of a woman richer than himself, but no particular woman corresponding to that notion. As Quine points out, the two readings are formally distinguished in some languages. Thus in French (as in other Romance languages — Quine's own example is in Spanish)¹ there are two ways of saying that John wants to marry a woman richer than himself, only one of which entails that *there is* a woman richer than himself whom John wants to marry (relational reading) :

Relational :

Jean veut épouser une femme qui *est* plus riche que lui [indicative]

Notional :

Jean veut épouser une femme qui *soit* plus riche que lui [subjunctive]

According to Quine, that sort of ambiguity also arises when we use a singular term, e.g. a name or a definite description, in the content clause of an attitude report. Thus a belief report such as 'John believes that the shortest spy is a spy' can be understood in two ways :

Relational reading

B³ (John, the shortest spy, z is a spy)

Notional reading

B² (John, the shortest spy is a spy)

On the first reading the sentence can be paraphrased as 'John believes of the shortest spy that he is a spy' ; not so on the second reading, where the description 'the shortest spy' contributes a notion (the concept of shortest spy) to the content of the ascribed belief. Relationally understood, the 'statement affirms a relation between [two] persons, and the persons remain so related independently of the names applied to them' (1961 : 142) ; so substitution of any other designation for the same person cannot affect truth-value. Indeed, on the relational reading of 'John believes that the shortest spy is a spy', where it means that John believes of the person who is in fact the shortest spy that he is a spy, if the shortest spy happens to be Jane's second son it can be inferred that John believes of Jane's second son that he is a spy. But if what John believes is the general proposition that (whoever is) the shortest spy is a spy, substitution does not preserve truth-value. Likewise, existential generalization is possible on the relational reading : if John believes of the shortest spy (a.k.a. Jane's second son) that he is a spy, then there is someone whom he believes to be a spy. Not so on the notional reading : John may believe that whoever is the shortest spy is a spy without believing of anyone in particular that he is a spy.

Quine applies the distinction to belief reports featuring a proper name, such as 'Tom believes that Cicero denounced Catiline'. That sentence too may be understood in two ways, he says : either transparently as saying of Cicero that Tom believes he denounced Catiline (relational reading), or opaquely as saying that Tom holds true the sentence 'Cicero denounced Catiline'.

¹ Quine's example is 'Procuro un perro que habla' vs 'Procuro un perro que hable' (Quine 1956 : 177).

Relational reading

B³ (Tom, Cicero, z *denounced Catiline*)

Notional reading

B² (Tom, *Cicero denounced Catiline*)

Again, substitution of identicals and existential generalization are only possible on the relational reading. That is the criterion Quine offers a for telling apart the two readings of attitude reports.

Interesting though it is, Quine's distinction between relational and notional attitude reports suffers from a significant difficulty ; a difficulty which Brian Loar detected in his first published paper, 'Reference and Propositional Attitudes' (1972). Quine, it seems, has conflated two different distinctions, which are both legitimate but should be kept apart (Recanati 2000a). One is the distinction between those attitudes which target particular individuals (e.g. John's desire to marry a particular woman richer than himself) and those which don't ; and the other is the distinction between transparent and opaque attitude reports. The problem for Quine is that singular attitudes targeting particular individuals can themselves be reported 'opaquely' as well as 'transparently'. If one assumes, as many people do, that to hold an attitude targeting a particular individual (as opposed to an attitude with general content) is to stand in a certain relation to that individual, then Quine's conflation of relationality and transparency becomes unacceptable ; for it entails that whenever an attitude report is opaque, the attitude it reports must belong to the nonrelational variety.²

2

Let us reconsider the distinction between the two types of attitude : those which involve a relation between the attitude holder and some particular object which the attitude is about, and those which do not involve such a relation. The former are sometimes called '*de re* attitudes'. The distinction is well illustrated by Quine's contrast between wanting a particular sloop and wanting mere 'relief from slooplessness'. The content of the desire is *singular* in one case (involving a particular sloop) and *general* in the other (wanting a sloop, but not anyone in particular). Likewise, 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 *individuals*: Quine and Carnap. Because the belief is irreducibly singular and concerns particular individuals, rather than whoever satisfies such and such conditions, having the belief involves being suitably related to the particular individuals in question ; and ascribing the belief involves affirming a relation between the believer and these individuals.

The idea that singular belief (belief about particular objects) is fundamentally relational has loomed large in anti-descriptivist theories of reference. Bach, for example, argues as follows :

² Thus Quine : 'If belief is taken transparently, then 'Tom believes that Cicero denounced Catiline' expresses an outright relation between the men Tom and Cicero... ; if belief is taken opaquely, then [that statement] expressly relates Tom to no man' (Quine 1960 : 145).

If *all* your thoughts about things could only be descriptive, your total conception of the world would be merely qualitative. You would never be related in thought to anything in particular. Thinking of something would never be a case of having it ‘in mind’, as we say colloquially, or as some philosophers have said, of being ‘*en rapport*’, in ‘cognitive contact’, or ‘epistemically intimate’ with it. But picturesque phrases aside, just what is this special relation? Whatever it is, it is different from that involved in thinking of something under a description. If we can even speak of a relation in the latter case, it is surely not a real (or natural) relation. Since the object of a descriptive thought is determined SATISFACTIONALLY, the fact that the thought is of that object does not require any connection between thought and object. However, the object of a *de re* thought is determined RELATIONALLY. For something to be the object of a *de re* thought, it must stand in a certain kind of relation to that very thought. (Bach 1987 : 12)

On this view singular attitudes (*de re* thoughts, in Bach’s terminology) have the singular content they have in virtue of the attitude holder’s standing in certain relations to the particular object or objects that feature in the content of the attitude. If the attitude holder was not suitably related to the objects in question, the thought would be purely qualitative and would not target any particular individual. Singular attitudes therefore are necessarily relational attitudes : whoever believes that Quine was a student of Carnap believes *of* Quine and Carnap that the former was a student of the latter. The relevant relations (between the attitude holder and the particular objects involved in the content of singular attitudes) descend from more basic, informational relations such as the relations of perceiving, of remembering or of hearing about. All these relations are genuine relations. If John perceives, remembers, or hears about the table, there is something which he sees, remembers or hears 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.)

By and large, Brian Loar accepted the distinction,³ even though, in his paper ‘The Semantics of Singular Terms’ (1976), he expressed dissatisfaction with the ‘radical two-use theory’ which I have just expounded and which Bach’s quotation illustrates. Loar’s dissatisfaction sprang from two sorts of consideration. First, he held that even in the relational (*de re*) case, the attitude holder is thinking of the referent through some ‘individual concept’ or mode of presentation ; and he took individual concepts to be generally expressible by means of definite descriptions. Second, he thought there was a continuum of cases rather than a sharp division between the two classes.⁴ These reservations led Loar to reformulate the distinction as a distinction between two types of individual concepts, those that are and those that are not ‘identifying’, where identifyingness is a graded property. Be that as it may, how exactly we cash out the distinction between the two types of attitude is not what matters for

³ ‘Some belief reports do imply that the believer stands in certain relations to particular items, and these relations are probably at least partially causal relations. Other belief reports do not have such import, but seem to assert something of a rather more general nature about the believer, as well as about his beliefs.’ (Loar 1972/2017 : 25)

⁴ ‘Consider this range of cases. Case (1) : we see a large empty shoe ; S says about its owner, whoever that may be, ‘He’s rather big’. (...) Case (2) : we see large footprints on the beach ; same utterance. Case (3) : we see a mound of sand, which we assume covers a man ; same utterance. Case (4) : like the last case, except that a leg is protruding. Case (5) : we see the man directly ; S says, pointing, ‘He’s rather big’.’ (Loar 1976/2017 : 51-52)

our present purposes. What matters is the relation between that distinction and *another* distinction, between transparent and opaque attitude ascriptions.

Quine more or less conflated the two distinctions : he held that ‘notional’ attitude reports are opaque, and permit neither substitution of identicals *salva veritate* nor existential generalization, while ‘relational’ attitude reports are transparent and allow for both. Thus, as we have seen, ‘John believes, of the shortest spy, that he is a spy’ entails that John believes, of Jane’s second son, that he is a spy (assuming the shortest spy = Jane’s second son), and also entails that there is someone of whom John believes that he is a spy. ‘The shortest spy’ occurs extensionally in such a statement, whose logical form is

B³ (John, the shortest spy, *z is a spy*)

In contrast, ‘the shortest spy’ occurs nonextensionally in ‘John believes that the shortest spy is a spy’ when this is understood notionally as ascribing to John a thought involving the *concept* of shortest spy. That is why Quine equated the two distinctions, and used nonextensionality as a criterion for nonrelationality. But that was a mistake, Loar pointed out. Some relational belief reports are transparent, but not all are.

Loar gives the following example of a belief report that is both relational and opaque :

(1) Michael believes that that masked man is a diplomat

Suppose the speaker utters (1) while pointing to a certain man (say, Brown) at a costume ball. The speaker thereby refers to Brown, and says of him that Michael believes him to be a diplomat. This looks like a relational belief report. Still, the report may well be opaque. It may be that Michael talked to the masked man during the ball, and inferred from his conversation with him that the man in question (whom he did not recognize as Brown) is a diplomat, even though Michael would strongly deny that Brown (whom he knows to be a dentist) is a diplomat. In this situation the report in (1) is both relational and opaque. It is relational because Michael is said to believe of a certain man, singled out by a pointing gesture, that he is a diplomat ; but it is opaque because, even though that masked man = Brown, substitution of ‘Brown’ for ‘that masked man’ turns (1), which is true, into (2), which is false in that situation :

(2) Michael believes that Brown is a diplomat

To be sure, (2) is true in that situation if it is understood ‘transparently’, that is, as merely saying that Michael believes of Brown (thought of in some way or other) that he is a diplomat. But (2) is false on the opaque reading, while (1) is true on the opaque reading. The point is precisely that (1) can be construed as opaque, and resisting substitution, while still being a relational belief report ascribing to Michael a belief about a particular individual (thought of in a particular way).

Because substitutivity fails, Quine would say that a belief report such as (1) is nonrelational ; for if it affirmed a relation between Michael and the masked man, that relation could be affirmed using any another designation for the latter, e.g. the name ‘Brown’. But Loar rebuts that argument. Substitution may be blocked, he says, because the singular term ‘that masked man’ plays a *dual role* ; it both refers to the person the belief is about (the masked man, a.k.a. Brown) while mirroring the way Michael thinks of him when he thinks he is a diplomat. Such dual role, he argues, is similar to the dual role of the name ‘Giorgione’ in Quine’s own example

Giorgione was so-called because of his size.

Even though ‘Giorgione’ refers to Giorgione in that statement, it cannot be replaced by another name for Giorgione, say ‘Barbarelli’, *salva veritate*. That is because the name ‘Giorgione’ does double duty in the sentence : it refers to Giorgione (something which ‘Barbarelli’ can also do) but at the same time it is the antecedent for the demonstrative adverb in ‘so-called’. Substitution (of ‘Barbarelli’ for ‘Giorgione’) changes the interpretation of the adverb and thus affects truth-value :

Barbarelli was so-called because of his size

Substitution holds, however, if we *fix* the interpretation of ‘so-called’ so as to make it context-independent. From

Giorgione was *called* ‘Giorgione’ because of his size

it does follow that

Barbarelli was called ‘Giorgione’ because of his size.

Likewise for (1) : substitutivity is restored if we *fix* the aspect of its interpretation which is context-sensitive and depends on the choice of the referring expression. From

Michael believes that that masked man, *thought of as that masked man*, is a diplomat

it does follow that

Michael believes that Brown, thought of as that masked man, is a diplomat

Accordingly, Loar suggests the following logical form for a report such as (1) which is both relational and opaque :

B^3 (Michael, that masked man, *z is that masked man and z is a diplomat*)

This displays the dual role played by the singular term ‘that masked man’ : it refers to the individual Michael is belief-related to, *and* specifies how Michael thinks of that individual when he holds the reported belief.⁵

To be sure, Quine would insist that such a belief report is not *purely* relational ; the role of the singular term, ‘that masked man’, is not purely to specify the individual Michael’s

⁵ Another interesting example Loar gives involves the specific indefinite ‘a certain’:

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 points out that that statement ‘will often be taken to imply more than’ the purely relational ‘there is a cabinet member *x* such that Ralph believes *x* is a spy’ (Loar 1972/2017: 22). On the most natural reading, Ralph ascribes spyhood to the guy *thought as a cabinet member*. The logical form is :

$(\exists x) (x \text{ is a cabinet member} \ \& \ B^3 (\text{Ralph, } x, \text{ *z is a cabinet member and z is a spy*}))$

belief is about — it has an additional role and contributes the *notion* of masked man to the content of the reported belief. What Quine means by ‘relational’ is, indeed, *purely relational*. Purely relational belief reports are subject to the principle of substitutivity, in contrast to those belief reports that are not purely relational. Verbal matters aside, however, the intuitive distinction between relational and nonrelational attitudes which Quine himself appeals to when he gives his initial examples (wanting a particular sloop vs wanting relief from slooplessness, etc.) has to be recognized as distinct from the theoretical distinction that is of interest to him, that between the attitude reports that are transparent and license free substitution of identicals and those that are not. An adequate and comprehensive theory should combine the two distinctions, instead of conflating them.

3

In Loar’s framework, a belief report such as ‘John believes that his successor is a linguist’ has (at least) three different readings, not two (Loar 1972/2017 : 23, fn6).⁶ First, it has a transparent reading, paraphrasable as

- (i) John believes of the person who is in fact his successor that she is a linguist.

Thus understood the statement does not entail that John realizes that the person he takes to be a linguist is his successor. In that interpretation the description ‘his successor’ occurs extensionally and supports both substitution and existential generalization. It is *the speaker* who chooses the description ‘his successor’ to refer to the person whom John believes to be a linguist. In other words, the description does not occur in the content of the belief that is ascribed to John. This reading is purely relational, and its logical form is

B³ (John, his successor, z is a linguist)

The statement also has a purely notional, attributive reading, where the belief ascribed to John is a general belief about whoever is his successor, rather than a singular belief about a particular person :

- (ii) John believes that whoever is/will be his successor is a linguist

In this case, the description does occur in the content of the belief that is ascribed to John. Neither substitution nor existential generalization are possible. The logical form of the statement on that interpretation is

B² (John, his successor is a linguist)

But there is also a third, intermediate reading, which may be the most natural for that sort of statement. On that reading, as on the first reading, John’s belief is about a particular person (so existential generalization is possible) : John believes of that particular person, say Mary, that she is a linguist. Yet, in contrast to the first reading, ‘his successor’ does not occur extensionally. As on the second reading, the description does feature in the content of the belief that is ascribed to John :

- (iii) John believes of his successor (Mary), *thought of as his successor*, that she is a

⁶ Loar himself says ‘at least’. In Recanati 1993 : 386-93 I distinguish six possible readings for such a sentence.

linguist.

Assuming Mary is Peter's lover, we cannot infer from 'John believes that his successor is a linguist', thus understood, that John believes that Peter's lover is a linguist (unless we understand this last report transparently). For John may not know that Mary is Peter's lover, and he may have positive reasons to deny that Peter's lover is a linguist. The third reading corresponds to the following logical form :

B³ (John, his successor, *z is his successor and z is a linguist*)

As we have seen (§1), Quine applies the relational/notional distinction indifferently to belief reports involving definite descriptions, such as 'John believes that his successor is a linguist', and to belief reports involving proper names, such as 'Tom believes that Cicero denounced Catiline' or 'John believes that Mary is a linguist'. But definite descriptions, although they can be used referentially, can also be used attributively ; and when they are they do not refer to any particular object, in the strong, relational sense of 'refer'. So there is a sharp contrast to be drawn between a relational reading of 'John believes that his successor is a linguist', such as the transparent reading (i), and a nonrelational reading, such as the attributive reading (ii). Proper names, however, differ from definite descriptions in being intrinsically referential (i.e. marked as such).⁷ So 'John believes that Mary is a linguist' only possesses *relational* readings : the logical form is triadic in all cases and the report states a relation between John, Mary, and what John believes about her. True, the report has two readings, depending on whether it is understood as transparent or opaque, but this is not a contrast between a relational and a nonrelational reading (as the contrast between (i) and (ii) is), but a contrast between two relational readings, representable as

B³ (John, Mary, *z is a linguist*)

and

B³ (John, Mary, *z is Mary [or : z is called 'Mary'] and z is a linguist*)

Note that the same contrast exists between the two relational readings of 'John believes that his successor is a linguist', namely (i) and (iii).

Following Loar, I conclude that we need a more complex picture than Quine's. We need both the distinction between relational and nonrelational attitude reports (based upon the distinction between singular or referential attitudes, on the one hand, and general or nonreferential attitudes, on the other) and the distinction, within the class of relational attitude reports, between those that are opaque and those that are transparent. Relational attitude reports do, and nonrelational attitude reports do not, support existential quantification ; but among relational attitude reports, only those that are transparent support substitution of identicals (Figure 1).

⁷ In Recanati 1993 I argue that genuine singular terms (proper names, indexicals etc.) convey a lexical feature REF which encodes referentiality. Definite descriptions can be used referentially but they are not lexically referential (they can be 'token-referential' but are not 'type-referential').

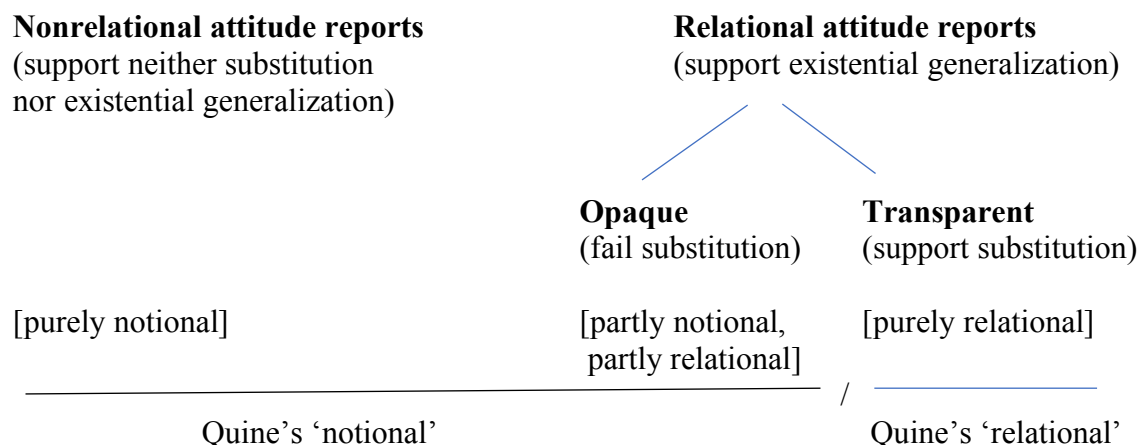


Figure 1 : Quine's picture amended

4

In relational belief reports, the speaker refers to the believer (say, Tom) as well as to another individual (say, Cicero), and ascribes to the former a belief about the latter – say the belief that he denounced Catiline. Now, all reference is under a ‘mode of presentation’. For example, Cicero can be referred to as ‘Cicero’, as ‘Tully’, or as the author of *De Senectute*. This applies not only to speech (talking about Cicero) but also to thought (thinking about Cicero). So, when a thought about Cicero is ascribed to Tom, there is a *double* reference to Cicero : the ascribed belief is about Cicero, and that means that the believer is taken to refer to Cicero in thought and to predicate of him the property of having denounced Catiline ; but the speaker who ascribes the belief herself refers to Cicero as the person the belief is about. Correspondingly, two modes of presentation are involved : the speaker presents Cicero in a certain way when she ascribes to Tom a belief about him ; and the believer, Tom, thinks of Cicero in a certain way when he judges that he denounced Catiline.⁸

In transparent belief reports, the way the believer thinks of the individual her belief is about is not specified : the report ascribes a belief concerning that individual, without specifying *how* the believer thinks of the individual in question. To be sure, the individual in question is presented in a certain way in the report – as Cicero, as Tully, or as the author of *De Senectute*. But this is irrelevant to the ascribed belief : the speaker chooses a particular way of referring to the individual the belief is about, but the mode of presentation thus selected is not ascribed to the believer. That is why substitution is possible : since the mode of presentation through which the speaker refers to the individual the belief is about is not ascribed to the believer, changing the mode of presentation does not affect the belief that is ascribed, hence the truth-conditions of the report. In contrast, opaque belief reports are such that the speaker's way of referring to the individual the belief is about is supposed to mirror the ascriber's own way of thinking of that individual. Opaquely understood, therefore, ‘Tom believes that Tully denounced Catiline’ is incorrect if Tom himself does not think of Cicero as ‘Tully’ when he predicates of him the property of having denounced Catiline.

Let me now introduce a bit of terminology. In *Direct Reference* (1993) I distinguished between the *exercised* mode of presentation and the *ascribed* mode of presentation. The

⁸ Of course, the same considerations which apply to Cicero also apply to Catiline : the speaker presents Catiline in a certain way when she ascribes to Tom a belief about him ; and the believer, Tom, thinks of Catiline in a certain way when he judges that Cicero denounced him.

ascribed mode of presentation is (unsurprisingly) the mode of presentation which is ascribed to the believer in opaque attitude ascriptions. The exercised mode of presentation is the mode of presentation through which the speaker refers to whatever she is referring to, e.g. the individual the ascribed belief is about. In terms of that distinction we can rephrase Loar's theory as follows. In opaque singular ascriptions, the same mode of presentation is *both* exercised by the speaker in referring to the individual the belief is about *and* ascribed to the believer as part of the content of her belief. That is why, on the opaque reading of 'John believes that the F is G' we find *two* occurrences of the description 'the F' in the corresponding logical form :

B³ (John, the F, *z is the F and z is G*)

The individual concept 'the F' serves both as the exercised mode of presentation and as the ascribed mode of presentation, in opaque belief reports. In transparent belief reports, the mode of presentation only plays the first role : that of exercised mode of presentation. It does not play the other role (since no mode of presentation is ascribed, in transparent belief reports).

This theory is a little too simple, as there are opaque singular reports where the exercised mode of presentation is *not* simultaneously ascribed to the believer. Before presenting such cases, however, I need to introduce another terminological distinction (also drawn from *Direct Reference*) between two types of mode of presentation : that between linguistic and psychological modes of presentation.

Psychological modes of presentation are modes of presentation occurring in thought – whether in the thought expressed by the speaker (and grasped by the understanding hearer), or in the thought ascribed to the believer. Both exercised and ascribed modes of presentation count as psychological modes of presentation. In contrast, the linguistic mode of presentation is the minimal mode of presentation of the referent conventionally encoded by referring expressions. This corresponds to what Loar calls the 'referential qualifier' (Loar 1976/2017 : 49). Thus a description 'the F' encodes the property of being F ; a pronoun such as 'she' encodes the property of being a female individual ; a name NN encodes the property of being called NN.

The linguistic mode of presentation is typically poorer and more schematic than the exercised mode of presentation, which fleshes it out. Thus consider the following well-known example from Brian Loar's paper on singular terms :

Suppose that Smith and Jones are unaware that the man being interviewed on television is someone they see on the train every morning and about whom, in that latter role, they have just been talking. Smith says 'He is a stockbroker', intending to refer to the man on television ; Jones takes Smith to be referring to the man on the train. Now Jones, as it happens, has correctly identified Smith's referent, since the man on television is the man on the train ; but he has failed to understand Smith's utterance. It would seem that, as Frege held, some 'manner of presentation' of the referent is, even on referential uses, essential to what is being communicated. (Loar 1976/2017 : 51)

Here the linguistic mode of presentation is the property of the referent which the pronoun 'he' encodes : the property of being a male individual, or possibly a *salient* male individual. But the psychological mode of presentation exercised by the speaker (and the comprehending hearer) is something much more specific like 'the man on television'. Grasping the exercised mode of presentation, on the basis of the linguistic mode of presentation and the context, is essential for understanding, Loar says.

It is tempting to construe the exercised mode of presentation as a ‘specification’ or ‘logical restriction’ of the linguistic mode of presentation, as the idea of ‘fleshing out’ suggests. Loar resisted that move, however, and I believe he was right. I follow him in holding that the linguistic mode of presentation *constrains* the exercised mode of presentation, without necessarily being like a skeleton which the exercised mode of presentation fleshes out. In my own story, the exercised mode of presentation is a *mental file* deployed in thinking about the referent, and what the referring expression encodes corresponds to a piece of information in the file. The speaker’s file and the hearer’s file are coordinated, inter alia, via the constraint that they both have to contain the piece of information encoded by the singular term (e.g. that the referent is called ‘Cicero’).⁹

5

What has been said thus far suggests the following picture:

Loar’s picture

In opaque belief ascriptions the mode of presentation exercised by the speaker in referring to the individual the ascribed belief is about does double duty : it is not only exercised by the speaker (and whoever properly understands the utterance) but also ascribed to the believer. So whichever relation the linguistic mode of presentation bears to the exercised mode of presentation (which it constrains), it also bears to the ascribed mode of presentation when the report is understood opaquely.

That picture is not general enough, however. Some opaque belief reports involve a single mode of presentation playing a double role, as in Loar’s analysis, but others don’t. Accordingly we need to distinguish between several notions, or two types, of opacity.

The notion of opacity we have been operating with so far is the following :

Opacity in the narrow sense

In opaque belief ascriptions the exercised mode of presentation, which the linguistic mode of presentation constrains, is simultaneously ascribed to the believer.

But opacity, in that sense, is not the opposite of the notion of ‘transparency’ which we characterized in section 4 :

Transparency

In transparent belief ascriptions, no mode of presentation is ascribed to the believer – only the referent (the object the belief is about) is specified.

⁹ Being mental particulars, mental files cannot be shared, strictly speaking. But the exercised mode of presentation is supposed to be *common* to all of those who properly understand the utterance (the speaker and his addressee, in a situation of successful communication). Likewise, opaque attitude reports have just been characterized as such that the mode of presentation exercised by the hearer is ascribed to the attitude holder : that means that the mode of presentation in question is shared by the speaker (and the understanding hearer) *and* the attitude holder. To make sense of these claims, ‘sharing’ has to be understood loosely. On that understanding, a mode of presentation is shared between individuals just in case the mental files they deploy have enough in common, i.e. just in case they all satisfy the constraint set up by the linguistic mode of presentation plus whatever constraints derive from the speaker’s communicative intention (in Loar’s example : the intention to speak of the man on television). The claim that, in opaque attitude reports (or some of them at least), the exercised mode of presentation and the ascribed mode of presentation are ‘identical’ will have to be understood in the same manner, i.e. as not requiring actual identity but only sufficient similarity.

The opposite of transparency, thus characterized, would correspond to a broader notion of opacity :

Opacity in the broad sense

In opaque belief ascriptions some way of thinking of the referent is ascribed to the believer.

The difference between opacity in the narrow sense and opacity in the broad sense is this : there is opacity in the broad sense provided some mode of presentation is ascribed to the believer, *even if that mode of presentation is distinct from the exercised mode of presentation*. There is opacity in the narrow sense only if the ascribed mode of presentation *is* the exercised mode of presentation.

Of course, the difference I have just mentioned will only materialize if there are attitude ascriptions such that (i) some mode of presentation is ascribed to the believer, while (ii) the mode of presentation exercised by the speaker is not ascribed to the believer. Are there such reports ? I think there are (and they are far from uncommon).¹⁰

Think of the following story. King Peter wants to know what his subjects sincerely think of him. So he disguises himself as a beggar and goes to the market place, only accompanied by one guard (also disguised). He mixes with the crowd and engages conversation whenever he can. At some point, he notices a woman who keeps looking at him suspiciously. As the woman turns away to talk to an acquaintance, the King whispers to the guard :

(3) She knows I am the King. We should leave.

Let's suppose the woman's name is 'Berta'. The propositional attitude which King Peter ascribes to Berta is singular : it's about King Peter himself. He says that she knows that *he* is the King. Is the ascription transparent ? Is Berta said to know, of King Peter, that he is the King ? Clearly not. Every subject in the Kingdom knows that King Peter is the King, and it follows that they all know of King Peter that he is the King. It is not that trivial bit of knowledge which (3) ascribes to Berta. What (3) ascribes to Berta a quite specific bit of knowledge : that (e.g.) *the short person looking like a beggar and located at such and such a place* is the King. In other words, Peter takes Berta to have seen through his disguise. The true belief he ascribes to her is opaque in the broad sense : it is a belief about Peter that involves a particular way of thinking of Peter. That way of thinking of Peter is ascribed to Berta. It is the 'ascribed mode of presentation'.

What about the exercised mode of presentation ? In ascribing to Berta a belief about himself, Peter refers to himself. How – under which mode of presentation – does he refer to himself ? He uses the first person : 'She knows that *I* am the King.' So he presents himself (the person Berta's belief is about) as the person making this utterance. This mode of presentation is the linguistic mode of presentation (Loar's referential qualifier) : 'I' encodes the property of being the speaker. That linguistic mode of presentation constrains the exercised mode of presentation, in the sense that whichever mental files the speaker (King Peter himself) and his hearer (the guard) deploy in thinking about the referent must contain the piece of information that the referent is currently making this utterance. That is, indeed, something of which both King Peter and the guard are aware : they both know that Peter is saying 'She knows that I am the King' to the guard. So the exercised mode of presentation is a way of thinking of Peter that somehow involves the awareness that Peter is making this

¹⁰ See Morgan 2011 : 177 for a detailed example.

utterance ; but Berta is unaware that Peter is making this utterance – she is looking away as Peter whispers aside to the guard. It follows that, in this example, the exercised mode of presentation and the ascribed mode of presentation come apart. The ascription is opaque in the broad sense (it is not transparent – a particular mode of presentation is ascribed to the believer) but it is not opaque in the narrow sense (the exercised mode of presentation is not ascribed to the believer).

I conclude that there are more cases of ‘relational opacity’ than those Loar can analyse by means of his ‘intermediate’ logical form for utterances of the type ‘*a* believes that the *F* is *G*’ :

$B_3(a, \text{the } F, \underline{z \text{ is the } F \text{ and } z \text{ is } G})$

The case in which there is coincidence between the mode of presentation exercised by the speaker in referring to what the belief is about and the mode of presentation ascribed to the believer as part of the content of belief is only a special case (opacity in the narrow sense). Another type of relational opacity involves two distinct modes of presentation, as per the following logical form :

$B_3(a, \text{the } F, \underline{z \text{ is the } H \text{ and } z \text{ is } G})$

Here the ascribed mode of presentation (‘the *H*’) and the exercised mode of presentation (‘the *F*’) come apart. That is what we found with example (3), which can be analysed roughly as

$K_3(\text{Berta, the speaker, } \underline{z \text{ is that person looking like a beggar and } z \text{ is the King}})$

6

In the ‘King’ example, the linguistic mode of presentation constrains the way the speaker and his addressee think of the referent, but it does not constrain the mode of presentation ascribed to the believer (Berta does not think of the referent as being the speaker). This suggests that, perhaps, the ascribed mode of presentation is constrained by the linguistic mode of presentation only when it is identical to the exercised mode of presentation, which is so constrained.

But is that true ? Or could there be cases in which the ascribed mode of presentation is constrained by the linguistic mode of presentation, even though the ascribed mode of presentation is not identical to the exercised mode of presentation ? To answer that question, we need to look at the notion of exercised mode of presentation more closely, in order to clarify it.

There are two things we know about the exercised mode of presentation. First the exercised mode of presentation is the way the speaker and/or the understanding hearer think(s) of the referent. The speaker and the hearer need not think of the referent in the same way – obviously, the speaker who says ‘I’ does not think of himself in the same way as the addressee — but at least their ways of thinking of the referent are coordinated. The coordination is secured, inter alia, via the linguistic mode of presentation, which constrains both the speaker’s way of thinking of the referent and the hearer’s. That is the second thing we know about the exercised mode of presentation : it is constrained by the linguistic mode of presentation.

These two things may seem to come apart in a class of cases with the following characteristics. The speaker and the hearer think of the referent in a certain way, while ascribing to the believer a different way of thinking of the referent. Thus far, this is like the

example discussed in the previous section. But here comes the difference : in the relevant class of cases, the linguistic mode of presentation constrains the believer's way of thinking of the referent, but not the speaker's or the hearer's. That is the opposite of the situation described in the last section.

Consider the following example. The speaker and the hearer know that Mary is not John's future bride ; but John thinks she is. In that situation the speaker can ascribe to John a singular belief about Mary, thought of as his future bride :

(4) John believes that his future bride is coming tonight

This is not to be understood as an attributive, purely notional ascription ('John believes that whoever will be his bride is coming tonight'), but as a *de re* ascription : John believes of Mary, thought of as his future bride, that she is coming tonight. This is another instance of opaque-relational report, but it cannot be ascribed the same logical form as Loar's example (1) ('John believes that that masked man is a diplomat') because the description 'his future bride' is not exercised by the speaker in referring to Mary (or so it seems). *The speaker does not think of Mary as John's future bride.* We can even imagine that he has compelling reasons to believe that John will actually marry Sofia, so for him the description 'John's future bride' refers to Sofia. (Let us additionally suppose that he is right and that John's future bride is Sofia.) Still, in uttering (4), he says that John believes of Mary, thought of as his future bride, that she is coming tonight. This cannot be rendered as

B₃ (John, his future bride, *z is his future bride and z is coming tonight*)

for the first occurrence of 'his future bride' would refer to Sofia ! What we need, rather, is something like

B₃ (John, Mary, *z is his future bride and z is coming tonight*)

Here, as in the 'King' example, the exercised mode of presentation and the ascribed mode of presentation come apart (or so it seems). But, in contrast to the 'King' example, it is the ascribed mode of presentation which the linguistic mode of presentation constrains. Here the property of the referent which is encoded by the singular term is the property of being John's future bride ; and it is John himself, not the speaker (nor the hearer), who thinks of Mary as his future bride. So it seems that we have found a class of cases in which the linguistic mode of presentation constrains the ascribed mode of presentation rather than the exercised mode of presentation.¹¹

But that may not be the right description of the case. It all depends on what we mean by 'exercised mode of presentation' (hence the need for clarification which I mentioned above). In standard cases the exercised mode of presentation is the way the speaker/hearer think of the referent, *and* it is constrained by the linguistic mode of presentation (which helps coordinate the speaker's and hearer's respective ways of thinking). In the situation I have described, however, the two things come apart : the linguistic mode of presentation encodes the property of being John's future bride but the speaker and the hearer do *not* think of the referent as being John's future bride. If we decide that the exercised mode of presentation is the way the speaker and the hearer actually think of the referent, we have to conclude, as I have tentatively done, that sometimes the exercised mode of presentation is not constrained

¹¹ Another example in the same category is mentioned by Jennifer Hornsby (who, like Loar, argues for the existence of reports that are both relational and opaque) : 'Bill thinks that the bank manager was rude, but it actually was a clerk about whom Bill reached that conclusion' (Hornsby 1977 : 34).

by the linguistic mode of presentation, while the ascribed mode of presentation is. But this conclusion is not forced upon us. We can also *stipulate* that the exercised mode of presentation has got to be constrained by the linguistic mode of presentation. After all, the exercised mode of presentation in a belief report is the mode of presentation one exercises in referring to the individual the belief is about, and one refers to the individual in question by using an expression which encodes certain of its alleged properties. On this view nothing is the exercised mode of presentation unless it is constrained by the linguistic mode of presentation – there is a constitutive link between the two notions. It follows that, even though the speaker and the hearer do not *think* of Mary as John’s future bride in the situation we have described, still in that context they do refer to Mary *as* John’s future bride. The mode of presentation they exercise in referring to Mary is not the way *they* think of Mary, but the way *John* thinks of Mary. In other words : the mode of presentation through which they refer to Mary is a *vicarious mode of presentation*, in such cases. They refer to Mary from John’s perspective, by putting themselves into John’s shoes. Pursuing this line of argument will make it possible to maintain that the ascribed mode of presentation is constrained by the linguistic mode of presentation only when it is identical to the exercised mode of presentation (which is always so constrained).

7

In several writings over the past forty years I have mentioned the important phenomenon of *oblique reference*. There is oblique reference to some individual *x* when the mode of presentation through which the speaker refers to *x* is not the way the speaker herself thinks of *x* but the way some other person thinks of *x*. In such cases, to understand the utterance you need to understand that the speaker is, as it were, *echoing* some other person and referring to *x* as that other person would.¹² That the referring expression is used echoically can be indicated by putting it within quotation marks, as in the following examples :

- (5) Hey, ‘your sister’ is coming over !
- (6) ‘Quine’ has still not finished his paper

In example (5), from Recanati 1986, the speaker refers to Ann, who is not the addressee’s sister (and who both the speaker and his addressee know not to be the addressee’s sister), but who is believed to be the addressee’s sister by a third party, James. The speaker is ironically mocking James by referring to Ann as the addressee’s sister, thus echoing James’ mistake. That is possible because James’s mistake is common knowledge between the speaker and his addressee. In that situation the speaker says of Ann, vicariously thought of as the addressee’s sister, that she is coming over. The hearer can respond, ‘No, she is not – she has stopped and is retracing her steps’, and himself refer to Ann by means of the pronoun ‘she’, which anaphorically picks up its reference from the echoic description. Similarly, in (6), from Recanati (2000b, 2001), the speaker ironically uses the name ‘Quine’ (in quotes) to refer to a philosopher, McPherson, whom a third party has misidentified as Quine a few days ago. In that context (6) says that McPherson has still not finished his paper. The speaker can respond, ‘He has – he told me over lunch that he’s just finished’ and himself refer to McPherson by means of the pronoun ‘he’, which anaphorically picks up its reference from the echoic name.

¹² The label ‘echoic’ comes from Sperber and Wilson (1981, 1986). Other labels are used for what is essentially the same phenomenon : ‘polyphony’ (Ducrot 1980a, 1980b, 1984, Recanati 1987), ‘staging’ (Clark 1996 : chapter 12) and ‘protagonist projection’ (Holton 1997, Stokke 2013).

There are at least two (compatible) ways to look at the phenomenon. From a strictly semantic point of view, we can account for the shift of reference of the description ‘your sister’ (which does not refer to the addressee’s sister, as it should, but to the person whom James takes to be the addressee’s sister) or of the name ‘Quine’ (which does not refer to Quine, as it should, but to the philosopher who was misidentified as Quine by the person the speaker is ironically mocking) by means of a *context-shifting operator* contributed by the (implicit or explicit) quotation marks around the expression. Thanks to the operator, the expression refers to what it would refer to in the shifted context (Recanati 1997, 2000b, 2008/2010). Or, in more cognitive terms, we can say that the linguistic mode of presentation constrains the exercised mode of presentation, as usual, and therefore imposes that the mental file deployed in thinking about the referent contain the piece of information encoded by the expression (that the referent is Quine — or is named ‘Quine’ — or that the referent is the addressee’s sister), but that, in this particular case, the mental file deployed in thinking about the referent is not the speaker’s (or hearer’s) ‘regular file’ about it but an ‘indexed file’, that is, a vicarious file indexed to some other person (and containing the mistaken piece of information).¹³

Oblique reference is particularly useful when one is describing a person’s intentional states or actions : it is natural to take that person’s perspective (to put oneself in their shoes) in specifying the object the state or action is directed toward. Thus consider the following passage from Ruth Millikan’s *Varieties of Meaning* :

When a kitten sees itself in the mirror, it tries to approach the other kitten, to smell it and touch it. Failing in this, it then tries to look behind the mirror. (Millikan 2004 : 123)

What the kitten sees in the mirror is what it tries (but fails) to approach, to smell and to touch. In referring to it, however, the speaker (Millikan) adopts two different perspectives : first, from her own, external perspective on the scene, she describes what the kitten sees as ‘itself’ ; then, in trying to make sense of the kitten’s actions, she adopts the kitten’s own perspective in describing what it is trying to approach, smell and touch : ‘the other kitten’. Of course there is no other kitten. Still, we don’t have to treat the description ‘the other kitten’ which Millikan uses as referentially vacuous. We can maintain that the description refers to what the kitten sees, namely itself, but through a vicarious mode of presentation corresponding to the kitten’s own perspective : when it looks at itself in the mirror, the kitten mistakes itself for another kitten. This perspective shift could be made explicit by putting the description within quotation marks :

When a kitten sees itself in the mirror, it tries to approach ‘the other kitten’, to smell it and touch it.

Or consider the following example, from Pinillos (2011), which I commented several times :

We were debating whether to investigate both Hesperus and Phosphorus ; but when we got evidence of their true identity, we immediately sent probes there.

In the first clause the speaker espouses the point of view of the unenlightened (herself and her peers before they came to realize that Hesperus is Phosphorus), and she refers to Venus

¹³ On indexed files see Recanati 2012 and 2016.

twice, via two distinct mental files (respectively associated with the names ‘Hesperus’ and ‘Phosphorus’) rather than via a single inclusive Venus-file corresponding to her current point of view. The inclusive Venus-file is associated with the demonstrative adverb ‘there’ at the end of the second clause. That second clause expresses the subject’s current point of view while the first clause is phrased from the point of view of the speaker and her peers before they learnt the identity (when they thought there were two distinct stars). The exercised modes of presentation associated respectively with each of the two singular terms ‘Hesperus’ and ‘Phosphorus’ in the first clause are mental files indexed to the speaker and her peers at the time of the deliberation which is reported (before they merged the two files into a single, inclusive file). In contrast, when she refers to Venus for the third time at the end of the second clause, the speaker deploys her regular, inclusive Venus-file. Again, we could make the perspective shift explicit by enclosing ‘Hesperus’ and ‘Phosphorus’ within quotation marks.

When a singular belief report is opaque in the broad sense, that is, when a particular mode of presentation of the individual the belief is about is ascribed to the believer, it is always possible to refer obliquely to that individual by vicariously deploying the very mode of presentation which is ascribed to the believer. Thus instead of ‘She knows that I am the King’, King Peter could have said : ‘She knows that *that funny beggar talking to you* is the King’. That possibility is exploited in the Pinillos example, and it is the key to understanding example (4), to which I finally return :

(4) John believes that his future bride is coming tonight

The description ‘his future bride’ can be construed as echoic, and thus as referring not to John’s future bride (Sofia) but, obliquely, to *the person whom John takes to be* his future bride (Mary). On that interpretation, the mode of presentation which the speaker (and his understanding hearer) exercise in thinking about Mary is not their ‘regular file’ about Mary, but a vicarious file indexed to John, and containing the mistaken piece of information that she is his future bride. The logical form of (4), thus understood, is

B³ (John, ‘his future bride’, *z is his future bride and z is coming tonight*)

where the quotation marks around the description filling the second argument slot of the triadic belief relation indicate that the description should be evaluated not with respect to the actual world but with respect to what John takes to the actual world (oblique reference). Since, in all the worlds compatible with John’s beliefs, his future bride is Mary, the echoic description ‘his future bride’ refers to Mary, and the report ascribes to John a singular belief about her.

In this type of case, the mode of presentation one exercises in referring obliquely to the object the belief is about is the mode of presentation we simultaneously ascribe to the believer. So, as in Loar’s example (‘that masked man’), the ascribed mode of presentation and the exercised mode of presentation are identical ; but in contrast to Loar’s example, it is the exercised mode of presentation which is parasitic on the ascribed mode of presentation (via the mechanism of indexed files and oblique reference). The difference can be presented as follows :

- In the ‘masked man’ type of case the mode of presentation exercised by the speaker in referring to the individual the ascribed belief is about does double duty : it is not only exercised by the speaker (and whoever properly understands the utterance) but also ascribed to the believer.

- In the ‘future bride’ type of case the mode of presentation ascribed to the believer does double duty : it is not only ascribed to the believer but also (vicariously) exercised by the speaker in obliquely referring to what the belief is about.

In both types of case, the ascribed mode of presentation and the exercised mode of presentation are identical. This is in contrast to the ‘King’ type of case, where the mode of presentation ascribed to the believer is distinct from the mode of presentation exercised by the speaker (and by whoever properly understands the utterance).

8

In the first half of this chapter I presented Loar’s criticism of Quine’s distinction between notional and relational attitude reports. I defended the claim that, whenever the content clause of an attitude report contains a genuine singular term referring to some individual the attitude is about, the report has to be construed as ‘relational’, even if it is opaque and fails substitutivity (Figure 1). I fully agree with Loar’s views on this matter, and have adopted his more comprehensive picture, based upon ‘a seemingly acceptable kind of statement of which Quine’s analysis fails to make sense, but which our expanded version accommodates without difficulty’ (Loar 1972/2017 : 22). In the second half of the chapter, I argued that Loar’s alternative picture is still not comprehensive enough, as there are more cases of relational opacity than it can accommodate. I therefore attempted to make the picture more complex. The following chart recapitulates some of the distinctions I have made, and the main examples I have used (Figure 2).

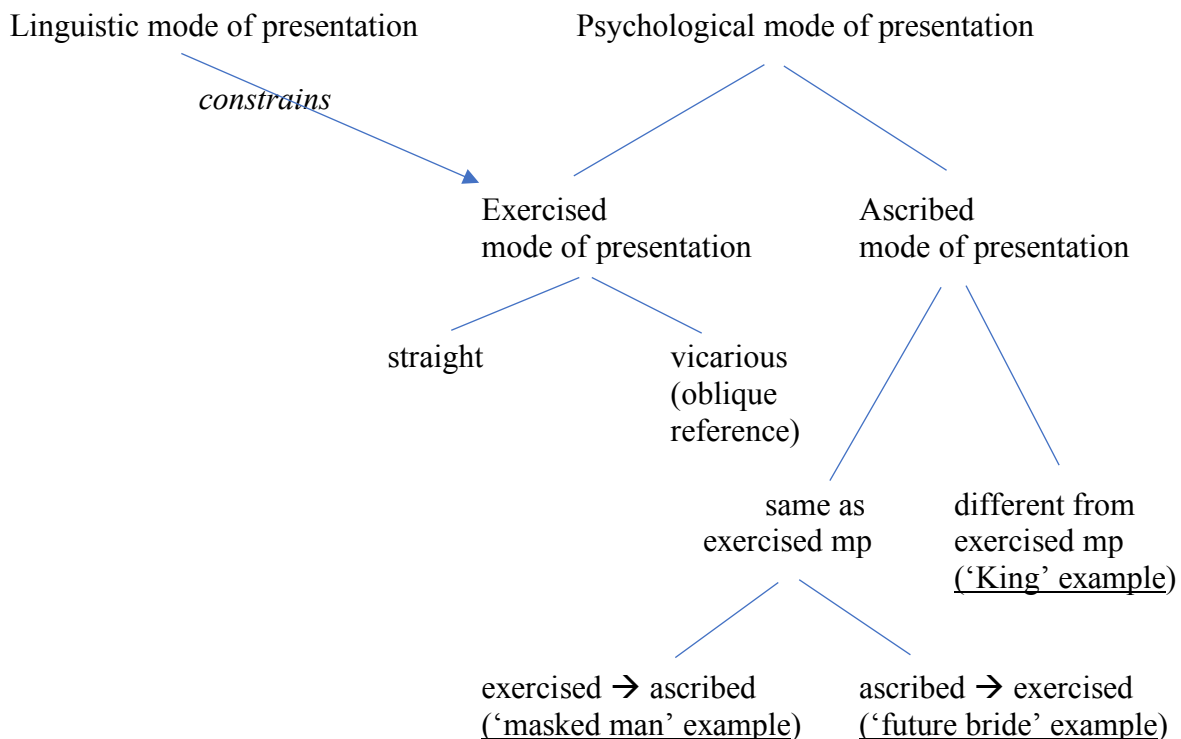


Figure 2 : Varieties of mode of presentation

References

- Bach, K. (1987) *Thought and Reference*. Oxford: Clarendon Press.
- Clark, H. (1996) *Using Language*. Cambridge : Cambridge University Press.
- Ducrot, O. (1980a) Analyses pragmatiques. *Communications* 32 : 11-60.
- Ducrot, O. (1980b) Analyse de texte et linguistique de l'énonciation. In O. Ducrot *et al. Les Mots du discours*, 7-56. Paris : Editions de Minuit.
- Ducrot, O. (1984) Esquisse d'une théorie polyphonique de l'énonciation. In O. Ducrot *Le Dire et le Dit*, 171-233. Paris : Editions de Minuit.
- Holton, R. (1997) Some Telling Examples: A Reply to Tsohatzidis. *Journal of Pragmatics* 28 : 625-628.
- Hornsby, J. (1977) Singular Terms in Contexts of Propositional Attitude. *Mind* 86: 31-48.
- Loar, B. (1972) Reference and Propositional Attitudes. *Philosophical Review* 81:43-62.
Reprinted in Loar 2017, 14-28.
- Loar, B. (1976) The Semantics of Singular Terms. *Philosophical Studies* 30: 353-77.
Reprinted in Loar 2017, 48-68.
- Loar, B. (2017) *Consciousness and Meaning : Selected Essays*. Oxford : Oxford University Press.
- Morgan, D. (2011) *First Person Thinking*. PhD dissertation, University of Oxford.
- Pinillos, A. (2011) Coreference and Meaning. *Philosophical Studies* 154: 301-24.
- Quine, W.V.O. (1956) Quantifiers and Propositional Attitudes. *Journal of Philosophy* 53 : 177-87.
- 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.
- Recanati, F. (1986) Contextual Dependence and Definite Descriptions. *Proceedings of the Aristotelian Society* 87 : 57-73.
- Recanati, F. (1987) *Meaning and Force*. Cambridge : Cambridge University Press.
- Recanati, F. (1993) *Direct Reference : From Language to Thought*. Oxford : Blackwell.
- Recanati, F. (1997) Can We Believe What We Do Not Understand ? *Mind and Language* 12 : 84-100.
- Recanati, F. (2000a) Opacity and the Attitudes. In Orenstein, A. and Kotatko, P. (eds) *Knowledge, Language and Logic : Questions for Quine*, 367-406. Dordrecht : Kluwer.
- Recanati, F. (2000b) *Oratio Obliqua, Oratio Recta*. Cambridge, Mass. : MIT Press/Bradford Books.
- Recanati, F. (2001) Open Quotation. *Mind* 110 : 637-87. (Revised version in Recanati 2010, chapter 7.)
- Recanati, F. (2008) Open Quotation Revisited. *Philosophical Perspectives* 22 : 443-71.
(Revised version in Recanati 2010, chapter 8.)
- Recanati, F. (2010) *Truth-Conditional Pragmatics*. Oxford : Oxford University Press.
- Recanati, F. (2012) *Mental Files*. Oxford : Oxford University Press.
- Recanati, F. (2016) *Mental Files in Flux*. Oxford : Oxford University Press.
- Sperber, D. and Wilson, D. (1981) Irony and the Use/Mention Distinction. In P. Cole (ed.) *Radical Pragmatics*, 295-318. London : Academic Press.
- Sperber, D. and Wilson, D. (1986) *Relevance*. Oxford : Blackwell.
- Stokke, A. (2013) Protagonist Projection. *Mind and Language* 28 : 204-32.