Metarepresentation in Linguistic Communication[*]

DEIRDRE WILSON

Abstract

This paper is designed to illustrate and consider the relations between three types of metarepresentational ability used in verbal comprehension: the ability to metarepresent attributed thoughts, the ability to metarepresent attributed utterances, and the ability to metarepresent abstract, non-attributed representations (e.g. sentence types, utterance types, propositions). Aspects of these abilities have been separately considered in the literatures on "theory of mind", Gricean pragmatics and quotation. The aim of this paper is to show how the results of these separate strands of research might be integrated with an empirically plausible pragmatic theory.

1 Introduction

Several strands of research on metarepresentation have a bearing on the study of linguistic communication. On the whole, there has been little interaction among them, and the possibility of integrating them with an empirically plausible pragmatic theory has not been much explored. This paper has two main aims: to illustrate the depth and variety of metarepresentational abilities deployed in linguistic communication, and to argue that a pragmatic account of these abilities can both benefit from and provide useful evidence for the study of more general metarepresentational abilities.

A metarepresentation is a representation of a representation: a higher-order representation with a lower-order representation embedded within it. The different strands of research on metarepresentation that have a bearing on the study of linguistic communication vary in the type of metarepresentations involved and the use to which they are put. First, there is the philosophical and psychological literature on mindreading (or "theory of mind"), which deals with the ability to form *thoughts* about *attributed thoughts* (Carruthers & Smith 1996; Davies & Stone 1995a, 1995b; Whiten 1991).

[*] I would like to thank Richard Breheny, Robyn Carston, Herb Clark, Steven Davis, Corinne Iten, Eun-Ju Noh, Milena Nuti, Anna Papafragou, François Recanati, Neil Smith, Dan Sperber, Tim Wharton and the members of CREA for interesting discussions on the topic of this paper. I would also like to thank CREA for its generous hospitality while the paper was being written.

Suppose a child sees a ball being put into a box. Having formed the thought in (1), he may go on, by observing his companions, to form thoughts of the type in (2):

(1) The ball is in the box.
(2) a. John thinks the ball is in the box.
    b. John thinks the ball is not in the box.
    c. John thinks Sue thinks the ball is in the box.
    d. John thinks Sue thinks the ball is not in the box.

There is a now a substantial body of work on how this metapsychological ability develops and how it may break down. It may be present to varying degrees. People may differ, for example, in their ability to attribute to others beliefs incompatible with their own. A child who believes (1) and lacks this ability would be limited to the metarepresentations in (2a) and (2c). A child with first-order "theory of mind" could attribute to others beliefs that differ from his own (as in (2b)); and one with second-order "theory of mind" could attribute to others beliefs about the beliefs of others which differ from his own (as in (2d)) (Astington, Harris & Olson 1988; Fodor 1992; Frye & Moore 1991; Gopnik & Wellman 1992; Leslie 1987; Lewis & Mitchell 1994; Scholl & Leslie 1999; Smith & Tsimpli 1995). Autistic people are typically said to be lacking in first- or second-order metapsychological abilities of this type (Baron-Cohen 1995; Baron-Cohen, Leslie & Frith 1985; Baron-Cohen, Tager-Flusberg & Cohen 1993; Happé 1993, 1994; Leslie 1991).

Second, there is the Gricean pragmatic literature on the attribution of speaker meanings. Grice shifted attention away from a code model of communication and towards an inferential account in which the formation and recognition of communicators' intentions was central. Thanks to his work, the idea that verbal comprehension is a form of mindreading has been relatively uncontroversial in pragmatics for more than thirty years (Bach & Harnish 1979; Davis 1991; Kasher 1998; Grice 1989; Levinson 1983; Neale 1992; Sperber & Wilson 1986/95). Grice treats the comprehension process as starting from a metarepresentation of an *attributed utterance* and ending with a metarepresentation of an *attributed thought*. Suppose Mary says (3) to Peter:

(3) You are neglecting your job.

In understanding her utterance, Peter might entertain a series of metarepresentations of the type in (4):

(4) a. Mary said, "You are neglecting your job."
b. Mary said that I am neglecting my job.
c. Mary believes that I am neglecting my job.
d. Mary intends me to believe that I am neglecting my job.
e. Mary intends me to believe that she intends me to believe that I am neglecting my job.

Unlike the literature on mindreading, the Gricean pragmatic literature deals with the specific metacommunicative ability to attribute speaker meanings on the basis of utterances. It might thus be seen as forming a bridge between the literature on mindreading and the philosophical, literary and linguistic literature on quotation, which is the third strand of research on metarepresentation that I will look at here.

The literature on quotation is mainly concerned with *utterances* about *attributed utterances*. Unlike the Gricean pragmatic literature, it deals with a type of metarepresentation used not *identifying* the speaker's meaning but as *part of* the speaker's meaning. For example, Peter might report Mary's utterance in (3) in one of the following ways:

(5) a. Mary said to me, "You are neglecting your job."
b. Mary told me I was not working hard enough.
c. According to Mary, I am "neglecting" my work.
d. Mary was pretty rude to me. I am neglecting my job!

(5) illustrates the four main types of quotation discussed in the literature: direct quotation, as in (5a), indirect quotation, as in (5b), mixed direct and indirect quotation, in (5c), and free indirect quotation, in (5d). Here, both the higher-order representation and the lower-order representations are utterances, and both are components of the speaker's meaning: they are part of what Peter intends to communicate by uttering (5a-d) (Cappelen & Lepore 1997; Coulmas 1986; Davidson 1968, 1979; McHale 1978; Noh 1998a; Partee 1973; Saka 1998).

So far, all the lower-order metarepresentations I have looked at have been attributed utterances or thoughts. There is a further, more disparate literature on non-attributive representations of a more abstract nature, linguistic, logical or conceptual. Consider the examples in (6):

(6) a. 'Dragonflies are beautiful' is a sentence of English.
b. 'Shut up' is rude.
c. It's true that tulips are flowers.
d. *Roses and daisies are flowers* entails that roses are flowers.
e. I like the name 'Petronella'.

f. 'Abeille' is not a word of English.
g. *Tulip* implies *flower*.

Here the higher-order representation is an *utterance* or *thought* and the lower-order representation is an *abstract representation*: for example, a sentence type, as in (6a), an utterance type, as in (6b), a proposition, as in (6c-d), a name, as in (6e), a word, as in (6f), or a concept, as in (6g). Such cases have been approached from a variety of perspectives: for example, the philosophical literature on quotation includes some discussion of non-attributive *mentions* of words or concepts (see also Garver 1965); and the ability to make grammaticality judgements, to think about sentence or utterance types, or to consider evidential or entailment relations among propositions or thought types, has given rise to a substantial experimental and developmental literature (Gombert 1990; Morris & Sloutsky 1998; Overton 1990).

Metarepresentation, then, involves a higher-order representation with a lower-order representation embedded inside it. The higher-order representation is generally an utterance or a thought. Three main types of lower-order representation have been investigated*: public representations*, e.g. utterances; *mental representations*, e.g. thoughts; and *abstract representations*, e.g. sentences, propositions. How do these metarepresentational abilities fit together, with each other and with the architecture of the mind? I will argue that it is worth considering them together and attempting to integrate them with an empirically plausible pragmatic theory. In the next section, I will consider how the Gricean metacommunicative ability used in attributing speaker meanings might fit with the more general metapsychological abilities studied in the literature on mindreading, and argue that some of Grice's assumptions about pragmatics must be modified if a serious attempt at integration is to be made. In the following section, I will sketch a pragmatic theory which might fit better with existing research on mindreading. In later sections, I will show how this theory might help with the analysis of quotation and other types of linguistic metarepresentation.

2 Gricean pragmatics and mindreading

Grice sees both communicator and audience as deeply involved in metarepresentation: the communicator in metarepresenting the thoughts she wants to convey, the audience in metarepresenting the communicator's intentions. Clearly, this metacommunicative ability has a lot in common with the more general mindreading ability illustrated in (2). It is conceivable that there is no difference between them, and that the ability to identify speaker meanings is nothing but the general mindreading ability applied to a specific, communicative domain. Arguably, this is the approach that Grice himself would have

favoured (see Sperber, forthcoming). An alternative hypothesis is that the metacommunicative ability is a specialisation of the more general mindreading ability, developed for use in the communicative domain. This possibility is currently being explored in work on relevance theory, and is the one I will develop here (Sperber 1996, forthcoming; Origgi & Sperber, forthcoming).

If either of these possibilities is to be seriously investigated, however, some of Grice's pragmatic assumptions will have to be dropped. His framework does not fit straightforwardly with existing research on mindreading, for several reasons. In the first place, his conception of communication involves, if anything, not too little metarepresentation but too much. For a Gricean speaker's meaning to be conveyed, the speaker's intentions must be not merely recognised but transparent, in a sense that seems to be definable only in terms of an infinite series of metarepresentations. The speaker must not only (a) intend to inform the hearer of something, and (b) intend the hearer to recognise this informative intention, but (c) intend the hearer to recognise the higher-order intention in (b), and so on ad infinitum. In other words, for a speaker's meaning to be conveyed, the speaker's informative intention – and every contextual assumption needed to identify it – must become *mutually known* (Recanati 1986; Schiffer 1972; Searle 1969; Smith 1982; Sperber & Wilson 1986/1995). However theoretically justified this conclusion, it creates a practical problem: it is hard to see how an infinite series of metarepresentations could ever be mentally represented. The search for a definition of speaker meaning that would simultaneously satisfy the theoretical requirement of transparency and the practical requirement of psychological plausibility was a major preoccupation of early inferential accounts (Clark & Carlson 1981; Clark & Marshall 1981; Garnham & Perner 1990; Gibbs 1987; Sperber & Wilson 1986, 1987, 1990a).

A second problem is that Grice seems to have thought of meaning attribution as involving a form of conscious, discursive reasoning quite unlike the spontaneous inferences deployed in mindreading. Here is his "working out schema" for the identification of conversational implicatures (Grice 1975, 50):

> *Grice's "working out schema" for conversational implicatures:*
> (a) He has said that *p*.
> (b) There is no reason to suppose that he is not observing the maxims, or at least the CP [= Co-operative Principle].
> (c) He could not be doing this unless he thought that *q*.
> (d) He knows (and knows that I know that he knows) that I can see that the supposition that he thinks that *q* is required.
> (e) He has done nothing to stop me thinking that *q*.
> (f) He intends me to think, or is at least willing to allow me to think, that *q*.
> (g) And so he has implicated that *q*.

It is hard to imagine even adults going through such lengthy chains of inference in the attribution of speaker meanings. Yet preverbal infants seem to be heavily involved in inferential communication, as the following example (from a 14-month-old infant) shows:

> Mother enters the room holding a cup of tea. Paul turns from his playpen in her direction and obviously sees it. (i) He cries vestigially and so attracts his mother's attention; immediately he points toward her and smacks his lips concurrently. [Paul's way of asking for food or drink.]
>
> *Mother*: No, you can't have this one, it's Andy's.
>
> Mother gives me [i.e. Andy Lock, the observer of the incident] the cup of tea, and I put it on the mantelpiece to cool. Paul crawls across to me and grasps my knees. (ii) I turn to look at him; he looks toward the mantelpiece and points, turns back to me, continues to point, and smacks his lips. (Lock 1980, 95-6).

Surveying a range of examples of this type, and noting the presence of such typical features of inferential communication as attracting the audience's attention, pointing, gaze alternation, and ritualised gestures such as lipsmacking and vestigial crying, Bretherton (1991, 57) concludes:

> I suggest that the most parsimonious explanation of these phenomena is that, by the end of the first year, infants have acquired a rudimentary ability to impute mental states to self and other … and, further, that they have begun to understand that one mind can be interfaced with another through conventional or mutually comprehensible signals.

While it is easy to accept that preverbal infants engage in inferential communication, it is hard to imagine them going through the sort of conscious, discursive reasoning illustrated in Grice's "working out schema".

In fact, the problem is more serious. Gricean pragmatics substantially underestimates the amount of inference involved in linguistic communication. As the "working out schema" shows, Grice saw the starting point for inferential comprehension as the recovery of a literal meaning (or "what is said"), which was determined independently of speakers' intentions. Yet there is good evidence that speakers' intentions help to determine not only what is implicated but also "what is said". This is most obvious in

disambiguation and reference resolution, but (as I will show in later sections) there is a very wide range of further cases in which sentence meaning substantially underdetermines "what is said" (Carston 1988, 1998; Sperber & Wilson 1986/1995, 1998b; Wilson & Sperber, in preparation). To accommodate these, Grice's "working out schema" for implicatures would have to be supplemented with further schemas designed to deal with disambiguation, reference assignment, and other linguistically underspecified aspects of "what is said": for example, resolution of lexically vague expressions, and interpretation of semantically incomplete expressions like 'too big'. While reflective inferences of this type do occur (for example in repairing misunderstandings or reading the later work of Henry James), disambiguation and reference assignment are in general intuitive processes which take place below the level of consciousness, and an adequate pragmatic theory should recognise this.

The ability to engage in inferential comprehension plays a role in language acquisition, where there is experimental evidence that young children attribute speaker intentions in acquiring lexical meanings. In one study (Tomasello & Kruger 1992), the experimenter used a novel verb in telling a child what she was about to do. She then performed an apparently accidental action (marked by saying "Whoops"), and an apparently intended action (marked by saying "There!" and looking pleased.). The child assumed that the verb described the apparently satisfactory action rather than the apparently accidental one. Bloom (1997: 10), who surveys a variety of examples of this type, concludes that "even very young children infer the referential intention of the speaker (through attention to cues that include line-of-regard and emotional indications of satisfaction) when determining the meaning of a new word." As with inferential communication in preverbal infants, it is easier to think of this as an intuitive rather than a reflective process.

If lexical comprehension involves an element of mindreading, the ability of autistic people to grasp an intended lexical meaning should also be impaired. Here is an illustration from the autobiography of someone with Asperger's syndrome:

> [During my first year at school], we were required to take naps each day. I vividly remember my teacher announcing, "Children, find your mats and take your nap." I refused. Again the teacher called my parents. Again my parents made their way to the school.
>
> 'Liane, why won't you take your nap?' my parents wondered of me.
>
> 'Because I can't.'
>
> 'You see!' the teacher said smugly.
>
> 'Why can't you take your nap?' my parents continued.
>
> 'Because I don't have a mat.'
>
> 'You most certainly do have a mat. There it is in your cubby,' the teacher replied.
>
> 'I do not have a mat.'
>
> 'You see what I mean?' the teacher asked my parents. 'She is an obstinate child.'
>
> 'Why do you say you don't have a mat?"' the folks asked, not giving up on me.
>
> 'That is not a mat. That is a rug,' I honestly and accurately replied.
>
> 'So it is,' said my father. 'Will you take a nap on your rug?'
>
> 'If she asks me to,' I said matter-of-factly
>
> I wasn't trying to be difficult, I was trying to do the right thing. The trouble was, the teacher assumed I understood language like other children. I did not. (Willey 1999, 19-20).

But the impairment in Asperger's syndrome seems to lie at the intuitive rather than the reflective level. If anything, failures at the intuitive level are compensated by an increase in the sort of reflective reasoning envisaged in Grice's "working out schema", as the following comment (from the same writer) suggests:

> If [my husband] were to tell me he was disappointed he had missed me at lunch, I would wonder if he meant to say he was sad – which is simply regretfully sorry; unhappy – which is somewhere between mad and sad; disheartened – which is a lonely sad; mad – which makes you want to argue with someone over what they had done; angry – which makes you want to ignore the person you are feeling this way towards; furious – which makes you want to spit; or none of the above. In order for me really to understand what people are saying I need much more than a few words mechanically placed together. (Willey 1999, 63)

Here the author describes a conscious attempt to resolve a lexical vagueness that most people would deal with spontaneously and unreflectively. This again suggests that the basic metacommunicative capacity is an intuitive rather than a reflective one.

Grice himself might not have been opposed to the idea of an intuitive metacommunicative capacity. What mattered to him was that this capacity – whether intuitive or reflective – was not code-based but inferential:

> The presence of a conversational implicature must be capable of being worked out; for even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if present at all) will not count as a CONVERSATIONAL implicature; it will be a CONVENTIONAL implicature. (Grice 1975, 50)

Grice's fundamental contribution to pragmatics was to show that much of verbal comprehension is inferential; but an empirically plausible pragmatic theory should also be concerned with how the inference processes go. Here, it is not the intuitions but the "working out schema" that need to be replaced.

A third problem is that, despite the elaborate-looking "working out schema", Grice's framework suggests no explicit procedure for identifying the content of particular speaker meanings. Grice showed that hearers have certain very general expectations – which he analysed in terms of a Co-operative Principle and maxims of truthfulness, informativeness, relevance and clarity – and look for meanings that satisfy those expectations. But how exactly is this done? How does the hearer decide, for instance, that someone who uses the word 'mat' intends to refer to a rug, or that someone who says "I was disappointed not to see you" is angry and expects an apology? If we look for guidance to the "working out schema" for implicatures, we find that the content of the implicature is introduced at step (c), but no explanation is given of how it is derived. In fact, the function of the "working out schema" is not to help the hearer construct a hypothesis about the content of the implicature, but merely to show how, once constructed, it might be confirmed as part of the speaker's meaning. But until we have some idea of how hypotheses about the speaker's meaning are constructed, we will be unable to see how the metacommunicative and metapsychological abilities might fit together.

In this paper, my main concern is with empirical questions about the role of metarepresentation in identifying the content of speaker meanings. In the next section, I will outline a pragmatic theory – relevance theory – which suggests a comprehension procedure that might replace Grice's "working out schema" and form the basis of a metacommunicative module. As to the theoretical problem of how to define speaker's meaning without getting into an infinite regress, Grice himself proposed a possible way out. He suggested that although full transparency in communication is not in practice

achievable (because of the infinity of metarepresentations required), communicators might simply *deem* it to be achieved (Grice 1982). This has the unfortunate consequence of making speaker's meaning an idealisation that is never achieved in real life. At the end of the next section, I will suggest an alternative solution which avoids this unfortunate consequence and shows how the theoretical goal of transparency and the practical goal of psychological plausibility might be reconciled.

3 Relevance theory and communication

Relevance theory (Sperber & Wilson 1986/1995, 1987) is based on a definition of relevance and two general principles: the *Cognitive Principle* that human cognition tends to be geared to the maximisation of relevance; and the *Communicative Principle* that utterances create expectations of relevance:

> *Cognitive principle of relevance:*
> Human cognition tends to be geared to the maximisation of relevance.
>
> *Communicative principle of relevance:*
> Every utterance (or other act of inferential communication) communicates a presumption of its own optimal relevance.

Relevance is treated as a property of inputs to cognitive processes and analysed in terms of the notions of cognitive effect and processing effort. When an input (for example, an utterance) is processed in a context of available assumptions, it may yield some cognitive effect (for example, by modifying or reorganising these assumptions). Other things being equal, the greater the cognitive effects, the greater the relevance of the input. However, the processing of the input, and the derivation of these effects, involves some mental effort. Other things being equal, the smaller the processing effort, the greater the relevance of the input.

It follows from the Cognitive Principle of Relevance that human attention and processing resources are allocated to information that seems relevant. It follows from the Communicative Principle of Relevance (and the definition of optimal relevance [Sperber & Wilson 1986/1995, 266-78]) that the speaker, by the very act of addressing someone, communicates that her utterance is the most relevant one compatible with her abilities and preferences, and is at least relevant enough to be worth his processing effort. This in turn suggests a comprehension procedure which might form the basis for a modularised metacommunicative ability.

Inferential comprehension starts from the recovery of a linguistically encoded sentence meaning, which is typically quite fragmentary and incomplete. The goal of pragmatic theory is to explain how the hearer, using available contextual information, develops this into a full-fledged speaker's meaning. The Communicative Principle of Relevance motivates the following comprehension procedure which, according to relevance theory, is automatically applied to the on-line processing of attended verbal inputs. The hearer takes the linguistically decoded sentence meaning; following a path of least effort in the accessing of contextual information, he enriches it at the explicit level and complements it at the implicit level, until the resulting interpretation meets his expectation of relevance; at which point, he stops:

> *Relevance-theoretic comprehension procedure*
> Follow a path of least effort in computing cognitive effects.
> (a) Consider interpretations in order of accessibility.
> (b) Stop when your expectation of relevance is satisfied.

The mutual adjustment of explicit content and implicatures, constrained by expectations of relevance, is the central feature of relevance-theoretic pragmatics (Carston 1998; Sperber & Wilson 1998b; Wilson & Sperber, in preparation).

The expectations of relevance created (and adjusted) in the course of the comprehension process may be more or less sophisticated. Sperber (1994) discusses three increasingly sophisticated strategies, each requiring an extra layer of metarepresentation, which might correspond to stages in pragmatic development. The simplest strategy is one of Naïve Optimism. A Naively Optimistic hearer looks for an interpretation that seems relevant enough: if he finds one, he assumes that it was the intended one and attributes it is a speaker's meaning; if he does not, he has no further resources, and communication will fail. In Sperber's terms, a Naively Optimistic hearer assumes that the speaker is both competent and benevolent: competent enough to avoid misunderstanding, and benevolent enough not to lead him astray. Suppose a mother tells her child:

(7) I'll write you a letter.

'Write you a letter' may mean *write a letter of the alphabet for you, write a message for you*, or *write a message to you*. The mother has spoken competently if the first interpretation that her son finds relevant enough is the intended one; she has spoken benevolently if this interpretation not only seems relevant but is genuinely so. A Naively Optimistic hearer has no need to think about the speaker's thoughts in identifying the speaker's meaning: the only time he needs to metarepresent the speaker's thoughts is when, having found an acceptable interpretation, he concludes that it is the intended one.

A more complex strategy, which requires an extra degree of metarepresentation, is one of Cautious Optimism. A Cautiously Optimistic hearer assumes that the speaker is benevolent, but not necessarily competent. Instead of taking the first interpretation he finds relevant enough and attributing it as the speaker's meaning, he can ask himself on what interpretation the speaker *might have thought* her utterance would be relevant enough. This extra layer of metarepresentation allows him to avoid misunderstanding in two types of case where a Naively Optimistic hearer would fail:

The first is the case of *accidental relevance*. An utterance is accidentally relevant when the first interpretation that seems relevant enough to the hearer is not the intended one. Suppose that – for reasons his mother could not plausibly have foreseen – the first interpretation of (7) that the child finds relevant enough is one on which his mother is offering to help him practise his handwriting. A Naively Optimistic hearer would accept this as the intended interpretation. A Cautiously Optimistic hearer would be able to consider whether his mother could have expected her utterance, on this interpretation, to be relevant enough to him.

An utterance may also be *accidentally irrelevant*. An obvious case is when someone mistakenly tells you something you already know. Another arises with slips of the tongue. Suppose Mary tells Peter:

(8) I've been feeding the penguins in Trafalgar Square.

A Naively Optimistic hearer would restrict himself to the linguistically encoded meaning, would be unable to find an acceptable interpretation, and communication would fail. By adopting a strategy of Cautious Optimism, and asking himself on what interpretation Mary *might have thought* her utterance would be relevant enough to him, Peter may conclude that she meant to say 'pigeon' instead of 'penguin'. Clearly, most ordinary hearers are capable of this.

While a Cautiously Optimistic hearer can deal with speaker incompetence, his assumption of speaker benevolence may still lead him astray. The strategy of Sophisticated Understanding allows hearers to cope with the fact that speakers are not always benevolent: they may intend an interpretation to *seem* relevant enough without in fact being so. For example, in saying (8), Mary may be lying about where she has been. A Cautiously Optimistic hearer might be able to cope with her slip of the tongue, but only if he does not realise she is lying: a benevolent communicator could not intend to inform him of something she knows to be false. Using the strategy of Sophisticated Understanding, Peter may be able to identify Mary's meaning even if he knows she is lying, by asking himself under what interpretation she *might have thought he would think* her utterance was relevant enough. In identifying the intended interpretation, he

therefore has to metarepresent Mary's thoughts about his thoughts. Most adult speakers are capable of this.

To sum up. A Naively Optimistic hearer need not metarepresent the speaker's thoughts at all in identifying the speaker's meaning: he simply takes the first interpretation that seems relevant enough and treats it as the intended one. A Cautiously Optimistic hearer considers what interpretation the speaker *might have thought* would be relevant enough: at the cost of an extra layer of metarepresentation, he can cope with cases where the speaker tries to be relevant enough, but fails. Finally, a hearer using the strategy of Sophisticated Understanding considers what interpretation the *speaker might have thought he would think* was relevant enough; at the cost of a further layer of metarepresentation, he can cope with deceptive cases in which nothing more than the appearance of relevance is attempted or achieved (see Sperber, forthcoming, for discussion).

These strategies have implications for the development of the metacommunicative ability. A child starting out as a Naive Optimist should make characteristic mistakes in comprehension (in disambiguation and reference assignment, for example), and there is some experimental evidence for this (Bezuidenhout and Sroda 1996, 1998). Roughly speaking, the move from Naive Optimism to Cautious Optimism coincides with the acquisition of first-order "theory of mind", and there should also be implications for verbal comprehension in people with autism and Asperger's syndrome (Leslie & Happé 1989; Happé 1993; for general discussion of the relation between relevance theory and mindreading, see Nuti, in preparation).

At the end of the last section, I pointed out an undesirable consequence of Grice's solution to the infinite-regress problem. On his account, transparency in communication, although *deemed* to be achieved, is never in fact achievable, so that full-fledged communication never occurs. Relevance theory suggests an alternative definition of communication that avoids this unfortunate consequence. The first step is to replace the notion of mutual knowledge (or mutual belief) with a notion of *mutual manifestness* (Sperber & Wilson 1986/95, chapter 1, section 8). Manifestness is a dispositional notion which is weaker than knowledge (or belief) in just the required way:

> *Manifestness*
> An assumption is manifest to an individual at a given time iff he is capable at that time of mentally representing it and accepting its representation as true or probably true.

An assumption cannot be known or believed without being explicitly represented;[1] but it can be manifest to an individual if it is merely capable of being non-demonstratively inferred. By defining communication in terms of a notion of mutual manifestness, the theoretical requirement of transparency and the practical requirement of psychological plausibility can be reconciled. (For discussion of mutual knowledge versus mutual manifestness, see Garnham and Perner 1990; Sperber & Wilson 1990a.)

Relevance theory analyses inferential communication in terms of two layers of intention: (a) the *informative intention* to make a certain set of assumptions manifest (or more manifest) to the audience, and (b) the *communicative intention* to make the informative intention mutually manifest (Sperber & Wilson 1986/95, chapter 1, sections 9-12):

> *Ostensive-inferential communication*
> The communicator produces a stimulus which makes it mutually manifest to communicator and audience that the communicator intends, by means of this stimulus, to make manifest or more manifest to the audience a set of assumptions **I.**

When the stimulus is an utterance, the content of the speaker's meaning is the set of assumptions **I** embedded under the informative intention. As long as the informative intention is made mutually manifest, transparency is achieved. An infinite series of metarepresentations is available in principle; however, it does not follow that each assumption in the series must be mentally represented. Which metarepresentations are actually constructed and processed in the course of interpreting a given utterance is an empirical question. On this account, the attribution of a full-fledged speaker's meaning involves a fourth-order metarepresentation of the type shown in (4e) above: she intends me to believe that she intends me to believe. This is complex enough to suggest a modularised metacommunicative ability, but finite enough to be implemented.

In this section, I have outlined a comprehension procedure which might form the basis for a modularised metacommunicative ability, itself a sub-part of the more general metapsychological ability, or "theory of mind". The procedure is governed by an expectation of relevance created and adjusted in the course of the comprehension process, which may be more or less sophisticated, with implications for development and breakdown. In the next section, I will turn to the content of the speaker's meaning

[1] Or at least deducible from assumptions explicitly represented. Since the full set of metarepresentations in a Gricean definition of speaker's meaning are not deducible from any finite subset, I will ignore this complication here. (See Sperber & Wilson 1990a for discussion).

(that is, the set of assumptions **I** embedded under the informative intention) and show that this may also contain a metarepresentational element which is very rich and varied.

4 Relevance theory and linguistic metarepresentation

4.1 Resemblance in linguistic metarepresentation

As noted above, the literature on quotation is mainly concerned with utterances about attributed utterances, such as those in (5) (repeated below):

(5) a. Mary said to me, "You are neglecting your job."
b. Mary told me I was not working hard enough.
c. According to Mary, I am “neglecting” my work.
d. Mary was pretty rude to me. I am neglecting my job!

Direct quotation, as in (5a), has been linked by different analysts to a variety of related phenomena: demonstrations, pretences, play-acting, mimesis, and non-serious actions. (Clark & Gerrig 1990; Recanati 1997; Sternberg 1982a; Walton 1990). When a literary example such as (9) is read out on the radio, it is easy to see why direct quotation has been treated as belonging to “a family of nonserious actions that includes practising, playing, acting and pretending” (Clark & Gerrig 1990, 766):

(9) “Out of the question”, says the coroner. “You have heard the boy. ‘Can't exactly say’ won't do, you know. We can't take that in a court of justice, gentlemen. It's terrible depravity. Put the boy aside.” (Dickens: *Bleak House*, chapter 11)

Similar claims have been made for free indirect quotation, as in (5d).

However, if we are interested in a notion of metarepresentation that extends to the full range of cases, public, mental and abstract, these analyses will not do. It is hard to see how notions such as pretence, mimesis and play-acting, which apply to public representations, can help with cases where the lower-order representation is chosen for its content or abstract properties: for example, the indirect speech report in (5b), the non-attributive mentions in (6), or indirect reports of thought such as (10):

(10) What, reduced to their simplest reciprocal form, were Bloom's thoughts about Stephen's thoughts about Bloom and Bloom's thoughts about Stephen's thoughts about Bloom's thoughts about Stephen? He thought that he thought that he was a jew, whereas he knew that he knew that he knew that he was not. (Joyce: *Ulysses* [Penguin, 602])

Nor do they help with cases where the higher-order representation is mental rather than public, as in the mental attributions of utterances or thoughts that underlie the metapsychological and metacommunicative abilities (cf. (2) and (4) above). What is worth retaining from these analyses is the idea that that quotation involves the exploitation of resemblances. I will argue that all varieties of metarepresentation, public, mental and abstract, can be analysed in terms of a notion of *representation by resemblance*, leaving the way open to a unified account.

In some of the literature on quotation, it has been assumed that identity rather than resemblance is the normal or typical case. Direct quotations are treated as verbatim reproductions of the original utterance, and indirect quotations as reproductions of its content. (For discussion, see Cappelen & Lepore 1997a, 1997b; Davidson 1968, 1979; Noh 1998a; Saka 1998). This assumption is too strong. In many cases, indirect quotation involves paraphrase, elaboration, or exaggeration rather than strict identity of content. For example, (5b) is a paraphrase of the original in (3), and it might be used to report a more remotely related utterance such as (11), which merely contextually implies or implicates that Peter is neglecting his job:

(11) You spend too much time at the theatre.

Particularly in academic circles, where even typographical errors are often reproduced verbatim, the idea that direct quotation is based on resemblance rather than identity may be harder to accept. But the degree of accuracy required in verbatim reporting depends on culture and circumstance (reproduction of phonetic features, hesitations, mispronunciations and repairs may or may not be relevant). Moreover, not all direct quotation is verbatim, as the following examples show:

(12) a. Descartes said, “I think, therefore I am.”
b. I looked at John and he's like, “What are you saying?”
c. And so the kid would say, “Blah blah blah?” [tentative voice with rising intonation] and his father would say “Blah blah blah” [in a strong blustery voice], and they would go on like that. (Clark and Gerrig 1990, 780)
d. And I said, “Well, it seemed to me to be an example of this this this this this this and this and this”, which it was you know. (Clark and Gerrig 1990, 780)

(12a) is a translation; the expression “he's like” in (12b) indicates that what follows should not be taken as a verbatim reproduction; and in (12c) and (12d) the expressions “blah blah blah” and “this this this” indicate very loose approximations indeed. (For discussion, see Clark & Gerrig 1990; Coulmas 1986; Gutt 1991; Wade & Clark 1993.)

A quotation, then, must merely resemble the original to some degree. Resemblance involves shared properties. As the above examples suggest, the resemblances may be of just any type: perceptual, linguistic, logical, mathematical, conceptual, sociolinguistic, stylistic, typographic. Typically, direct quotation, as in (5a), increases the salience of formal or linguistic properties, and indirect quotation, as in (5b), increases the salience of semantic or logical properties. We might call these resemblances *metalinguistic*, on the one hand, and *interpretive*, on the other (Sperber & Wilson 1986/95, chap 4, sections 7-9; Noh 1998a; Wilson & Sperber 1988b). Mixed quotation, as in (5c), exploits both metalinguistic and interpretive resemblances, while reports of thought, and metarepresentations of thought in general, are typically interpretive.

Interpretive resemblance is resemblance in content: that is, sharing of implications. Two representations resemble each other (in a context) to the extent that they share logical and contextual implications. The more implications they have in common, the more they resemble each other. Identity is a special case of resemblance, in which two representations share all their implications in every context. In interpreting a quotation, or more generally a linguistic metarepresentation, the hearer must make some assumption about the type and degree of resemblance involved. According to the relevance-theoretic comprehension procedure, he should not expect strict identity between representation and original: following a path of least effort, he should start with the most salient hypothesis about the intended resemblances, compute enough implications to satisfy his expectation of relevance, then stop. Resemblance, rather than identity, is the normal or typical case (Gibbs 1994; Sperber & Wilson 1990b, 1998; Wilson & Sperber, in preparation).

Developmental studies of quotation have provided useful data on the production side, tracing the development of propositional-attitude verbs such as ‘think’, ‘want’, ‘hope’, ‘fear’, for example (Bartsch & Wellman 1995; Bretherton & Beeghly 1982; Wellman 1990). Here I will look mainly at the comprehension side, and argue that language contains a huge variety of metarepresentational devices whose comprehension might interact in interesting ways with the metapsychological and metacommunicative abilities. I will also try to show that the recognition and interpretation of linguistic metarepresentations involves a substantial amount of pragmatic inference, bearing out my claim in the last section that Gricean pragmatics has considerably underestimated the inferential element in comprehension.

4.2 Decoding and inference in linguistic metarepresentation

The semantic and philosophical literature has been mainly concerned with overtly marked quotations such as (5a-c), whose presence is linguistically indicated by use of higher-order conceptual representations, e.g. “Mary said”, “Peter thought”. Literary and stylistic research has been more concerned with free indirect cases such as (5d), where

the presence, source and type of the metarepresentation is left to the reader to infer. Consider (13):

(13) Frederick reproached Elizabeth. She had behaved inconsiderately.

The second part of (13) has three possible interpretations: it may be an assertion by the narrator that Elizabeth had behaved inconsiderately, a free indirect report of what Frederick said, or a free indirect report of what he thought. The literature on "point of view" in fiction provides a wealth of information about clues to the presence of free indirect reporting, and critical procedures by which indeterminacies as to source and type of metarepresentation might be resolved (Banfield 1982; Cohn 1978; Fludernik 1993; Sternberg 1982b; Walton 1976).

To take just one example, consider (14), a passage from *Persuasion* which describes the reactions of the hero, Wentworth, on seeing the heroine, Anne Elliot, after a gap of many years. Anne's sister has just told her that Wentworth has informed a friend that Anne was "so altered he would not have known her". Anne is shocked and upset. Jane Austen continues:

(14) Frederick Wentworth had used such words, or something like them, but without an idea that they would be carried round to her. He had thought her wretchedly altered, and, in the first moment of appeal, had spoken as he felt. *He had not forgiven Anne Elliot. She had used him ill; deserted and disappointed him; and worse, she had shewn a feebleness of character in doing so, which his own decided, confident temper could not endure. She had given him up to oblige others. It had been the effect of over-persuasion. It had been weakness and timidity.*

He had been most warmly attached to her, and had never seen a woman since whom he thought her equal; but, except from some natural sensation of curiosity, he had no desire of meeting her again. Her power with him was gone for ever. (chapter 7)

As noted by Leech and Short (1989: 339) three different interpretations have been proposed for the italicised passage in (14). Mary Lascelles (1939/1965: 204) treats the first part as a straightforward authorial description, and the last part as a free indirect report of what Wentworth said. Wayne Booth (1961: 252) reads the whole passage as a free indirect report of Wentworth's thoughts. This disagreement has critical consequences. For Lascelles, the passage amounts to an "oversight" on Austen's part, since it fails to present events from the point of view of Anne Elliot, which is consistently maintained in the rest of the novel. For Booth, it is not an oversight at all.

By showing us Wentworth's thoughts at this one decisive point in the story, Austen creates a genuine doubt in our minds about what the outcome will be. “It is deliberate manipulation of inside views in order to destroy our conventional security. We are thus made to go along with Anne in her long and painful road to the discovery that Frederick loves her.” (Booth 1961: 252). Later critics have tended to prefer Booth's interpretation as yielding a more “coherent” reading.

In literary examples of this type, the interpretation process may be deliberate and time-consuming, calling on evidence from sources beyond the immediate context. In other cases, the presence of a quotation may be straightforwardly detected even though it is not overtly marked. Here are some examples from the “Question and Answer” column in a newspaper:

(15) a. Why is it that we curry favour?
b. Why is it that someone who tries to convert others proselytises?
c. Why is it that we trip the light fantastic if we go out for a good evening?
d. Why is it that we have to take off our shoes before entering a mosque?
e. Why is it that gorillas beat their chests?
f. Why is it that we get butterflies in our stomachs when we are nervous?

Although none of these questions contained quotation marks, some of them were clearly metalinguistic (“Why is it that we *say* we “curry favour”), while others were straightforwardly descriptive. In the published responses, (15a-c) were treated as metalinguistic, (15d-e) were treated as descriptive, and (17f) was treated as both. Intuitively, considerations of relevance help the reader decide how these questions were intended: it is easier to see how (15a-c) would be relevant as metalinguistic rather than descriptive questions, while the reverse is true for (15d-e). The relevance-theoretic comprehension procedure should shed light on how these utterances are understood.

A hearer following the relevance-theoretic comprehension procedure should consider interpretive hypotheses in order of accessibility. Having found an interpretation that satisfies his expectation of relevance, he should stop. The task of the speaker is to make the intended interpretation accessible enough to be picked out. Notice that the best way of doing this is not always to spell it out in full. In appropriate circumstances, the hearer may be able to infer some aspect of the intended interpretation with less effort than would be needed to decode it from a fully explicit prompt. Returning to (10), for example, it is relatively easy for the reader to infer that the pronouns in the final sentence must be understood as follows:

(10’) Bloom thought that Stephen thought that Bloom was a jew, whereas Bloom knew that Stephen knew that Bloom knew that Stephen was not.

This interpretation is justified in terms of both effort and effect. It is the most accessible one, since it is exactly patterned on the immediately preceding utterance, in which the intended referents are overtly marked. It is acceptable on the effect side, since it answers a question raised by the immediately preceding utterance, and achieves relevance thereby. The less explicit formulation in (10) is thus stylistically preferable to the one in (10'), which would cost the hearer some unnecessary linguistic effort.

Linguistic metarepresentations vary from the fully explicit and conceptual, as in (5a-b), to the fully inferred, as in (13). Most languages also have a range of quotative devices which indicate an attributive intention without foregrounding it to the degree shown in (5a-b). English has hearsay adverbs ('allegedly', 'reportedly'), adjectives ('self-confessed', 'so-called'), particles ('quote-unquote'), parentheticals ('as Chomsky says', 'according to Bill'), and noun-phrases ('Derrida's claim that', 'the suspect's allegation that'). French also has hearsay prepositions ('selon'), connectives ('puisque') and morphology (the "reportative conditional"); German has hearsay modals ('will'). Japanese has a hearsay particle ('tte') which, if added to the second part of (13), would mark it unambiguously as an attributed utterance; Sissala has an interpretive particle ('re') which does not distinguish between attributed utterances and thoughts. Inverted commas, "finger dancing" and intonation provide further orthographic and paralinguistic resources for indicating attributive use. These devices work in very varied ways, and their semantic properties and pragmatic effects deserve more attention than I can give them here. (For discussion, see Blass 1989, 1990; Ducrot 1983; Ifantidou-Trouki 1993; Ifantidou 1994; Itani 1996; Noh 1998a; Wilson & Sperber 1993).

Most languages also have a range of what might be thought of as self-quotative or self-attributive expressions, which add a further layer of metarepresentation to the communicated content. Parallel to "he thinks" and "he says" are "I think" and "I say"; and most of the hearsay expressions mentioned above have epistemic or illocutionary counterparts. Consider (16)-(17):

(16) a. Allegedly, the Health Service is on its last legs.
b. Confidentially, the Health Service is on its last legs.
c. Unfortunately, the Health Service is on its last legs.

(17) a. There will be riots, the security forces warn us.
b. There will be riots, I warn you.
c. There will be riots, I fear.

In (16a) and (17a), the parenthetical comment is used to attribute an utterance to someone other than the speaker; in (16b-c) and (17b-c), it carries speech-act or

propositional-attitude information about the speaker's own utterance (Blakemore 1992; Recanati 1987; Urmson 1963). Into this category of epistemic or illocutionary expressions fall mood indicators (declarative, imperative), evidentials ('doubtless'), attitudinal particles ('alas') and illocutionary-force indicators ('please'), which, by adding a higher-order metarepresentation to the basic layer of communicated content, might be seen as bridging the gap between the metacommunicative ability studied in Gricean pragmatics and the literature on quotation proper (Chafe & Nichols 1986; Clark 1991; Fillmore 1990; Ifantidou 1994; Papafragou 1998a, 1998b, 1998c; Recanati 1987; Wilson & Sperber 1988a, 1993).

As with freer forms of quotation, these higher-order metarepresentations need not be linguistically marked. Compare (18a) and (18b):

(18) a. The grass is wet, because it's raining.
b. It's raining, because the grass is wet.

Although syntactically similar, these utterances would normally be understood in different ways. (18a) would be understood as making the purely descriptive claim that the rain has caused the grass to get wet. The speaker of (18b) would normally be understood as communicating that the fact that the grass is wet has caused her to *say*, or *believe*, that it's raining. In (18a), the causal relation is between two states of affairs; in (18b), it is between a state of affairs and an utterance or thought. In interpreting (18b), the hearer must construct a higher-order representation of the type "she says", or "she thinks", and attribute it as part of the speaker's meaning (Noh 1998b; Papafragou 1998a, 1998b; Sweetser 1990; for developmental studies, see Noveck, Ho & Sera 1996; Papafragou 1998c).

In (18b), the inferred higher-order representation may be either epistemic or illocutionary. In other cases, this indeterminacy may be pragmatically resolved. Suppose someone comes up to me in the street and says (19):

(19) Your name is Deirdre Wilson.

The information that my name is Deirdre Wilson is patently irrelevant to me. What the speaker intends to communicate must be that she *knows*, or *believes*, that my name is Deirdre Wilson; only on this interpretation will (19) be relevant enough.

In still further cases, what has to be pragmatically resolved is whether some higher-order information made manifest by the utterance is part of the speaker's meaning or merely accidentally transmitted. Consider (20):

(20) a. *Mary* (whispering): I'm about to resign.
b. *Mary* (frowning): You're late.

c. *Mary* (puzzled): The radio's not working.

Here, paralinguistic features such as facial expression, gestures and intonation provide a clue to Mary's attitude to the proposition she is expressing, which may or may not be salient enough, and relevant enough, to be picked out by the relevance-theoretic comprehension procedure.[2] In the next section, I will look at a range of cases in which the speaker's attitude to an attributed utterance or thought makes a major contribution to relevance, and must be treated as part of the communicated content.

4.3 Reporting and echoing

The literature on quotation has been much concerned with reports of speech and thought, which achieve relevance mainly by informing the hearer about the content of the original. There is a wide range of further, *echoic*, cases which achieve relevance mainly by conveying the speaker's attitude to an attributed utterance or thought. Echoic utterances add an extra layer of metarepresentation to the communicated content, since not only the attribution but also the speaker's attitude must be represented.

The attitudes conveyed by echoic utterances are very rich and varied: the speaker may indicate that she agrees or disagrees with the original, is puzzled, angry, amused, intrigued, sceptical, etc., or any combination of these. Here I will limit myself to three broad types of attitude: endorsing, questioning and dissociative. Suppose Peter and Mary have been to see a film. As they come out, one of the following exchanges occurs:

(21) *Peter*: That was a fantastic film.
(22) *Mary*: a. [happily] Fantastic.
b. [puzzled] Fantastic?
c. [scornfully] Fantastic!

In (22a), Mary echoes Peter's utterance while indicating that she agrees with it; in (22b), she indicates that she is wondering about it; and in (22c) she indicates that she disagrees with it. The resulting interpretations might be as in (23):

(23) a. She believes I was right to say/think P.

[2] These are perhaps the clearest cases in which the mindreading ability makes a direct contribution to communicated content, by providing access to information about the speaker's mental states which may then be picked out by the relevance-theoretic comprehension procedure for attribution as a speaker's meaning.

b. She is wondering whether I was right to say/think P.
c. She believes I was wrong to say/think P.

Like regular quotations, echoic utterances may be metalinguistic or interpretive: the attitude expressed may be to the form of the original (e.g. a word, an accent, a pronunciation) or to its content. In (22b), for example, Mary may be wondering whether Peter meant to say the word 'fantastic', or to pronounce it as he did; or she may be wondering whether he really believes the film was fantastic, and why.

As with regular quotations, the speaker's attitude may be more or less overtly marked ("I agree that", "I doubt that", "I wonder whether"), or left to the hearer to infer, as in (22). Apart from intonation, facial expressions and other paralinguistic features, most languages also have various attitudinal devices, parallel to the hearsay devices above, which may increase the salience of the intended interpretation. Credal attitudes to attributed contents are conveyed by factive verbs ("he knows", "he admits", "they point out") and parentheticals ("as Chomsky says", "as these arguments have shown"). (On factives and credal attitudes, see Kiparsky & Kiparsky 1971; Sperber 1997). Questioning attitudes are conveyed by expressions such as "eh?", "right?", "aren't you?", as in (24):

(24) a. You're leaving, eh?
b. You don't want that piece of cake, right?
c. You're thinking of resigning, aren't you?

There is also a range of more or less colloquial dissociative expressions. Suppose Peter tells Mary that he's planning to enter the New York marathon, and she replies as in (25):

(25) a. You're bound to win, *I don't think.*
b. You're sure to win. *Not.*
c. You're going to run the marathon, *huh*!

Here, the addition of the italicised expressions makes it clear that the main clause is attributive, and that Mary's attitude is a sceptical or dissociative one.

A central claim of relevance theory has been that verbal irony is tacitly dissociative: the speaker expresses a wry, or sceptical, or mocking attitude to an attributed utterance or thought (Sperber & Wilson 1981, 1986/95, 1990b, 1998a; Wilson & Sperber 1992). Consider Mary's utterance in (22c) above. This is clearly both ironical and echoic. Relevance theory claims that it is ironical *because* it is echoic: irony consists in echoing a tacitly attributed thought or utterance with a tacitly dissociative attitude. This analysis has been experimentally tested, and the theoretical claims behind it have been much discussed (Clark and Gerrig 1984; Curcó 1998; Gibbs 1994; Jorgensen, Miller & Sperber 1984; Kreuz & Glucksberg 1989; Kumon-Nakamura, Glucksberg & Brown

1995; Martin 1992; Sperber 1984; Sperber & Wilson 1998a). There is good evidence that irony involves attributive metarepresentation, and that this extra layer of metarepresentation makes irony harder than metaphor to understand for people with autism who have not attained a second-order "theory of mind" (Happé 1993; on the development of metaphor and irony, see Winner 1988).

Verbal irony is interpretive: the speaker conveys a dissociative attitude to an attributed content. Parody might be thought of as its metalinguistic counterpart: the speaker conveys a dissociative attitude not (only) to an attributed content but to the style or form of the original. Typically, the resemblance is quite loose. Consider (26a), a mocking inversion of the saying in (26b):

(26) a. Our friends are always there when they need us.
b. Our friends are always there when we need them.

This is a case of echoic allusion, which allows the speaker to make a serious assertion with (26a) while simultaneously making fun of the related utterance in (26b).[3] A further type of case which is not normally treated as echoic is (27b):

(27) a. *Prince Charles:* Hello, I'm Prince Charles.
b. *Telephone operator:* And I'm the Queen of Sheba.

It has been suggested (I think by Dan Sperber) that the response in (27b) might be treated as an echoic allusion. The speaker aims to make salient a rather abstract property that her utterance shares with (27a): obvious falsehood or absurdity. In all these cases, the claim that attribution is based on resemblance rather than identity plays an important role.

Consider now the contrast between denial and negation. Negation is properly semantic; denial (typically conveyed by use of negative sentences) is a speech act, whose function is to reject some aspect of an attributed utterance or thought. In other words, denial is echoic. Here are some examples:

(28) a. *Peter*: Oh, you're in a miserable foul mood tonight.
b. *Mary*: I'm not in a miserable foul mood; I'm a little tired and would like to be left alone. (Carston 1996, 322)

[3] See Martin 1992 and Sperber & Wilson 1998a for discussion of such cases which, like many of my previous examples, present serious problems for traditional (non-attributive) analyses of verbal irony. For more general discussion of humour within a relevance-theoretic framework, see Curco 1995.

(29) Around here we don't eat tom[eiDouz] and we don't get stressed out. We eat tom[a:touz] and we get a little tense now and then. (Carston 1996, 320)

(30) Mozart's sonatas weren't for violin and piano, they were for piano and violin. (Horn 1989, 373)

(31) I didn't manage to trap two mongeese: I managed to trap two mongooses. (Horn 1989, 373)

In (28b), Mary echoes and rejects Peter's description of her; in (29), the speaker objects to the American pronunciation of 'tomatoes' and the expression 'stressed out'; in (30), she rejects the description of Mozart's sonatas as 'for violin and piano'; and in (31), she rejects the claim that she managed to trap two 'mongeese' rather than two 'mongooses'. Such denials fit straightforwardly into the pattern of previous sections. Like regular quotations, they may be interpretive, as in (28), or metalinguistic, as in (29)-(31). As with irony and free indirect quotations, the presence of the attributive element is not overtly marked.

In fact, the picture of denial just given is not the standard one. Linguists generally define denial as involving the rejection of an attributed *utterance*, treating rejections of attributed thoughts as cases of regular negation. For example, van der Sandt (1991: 331) claims that the "essential function" of echoic denials is "to object to a previous utterance". His category of denials would include (28)-(31), which all metarepresent attributed utterances, but would exclude rejections of attributed thoughts. Horn (1989: 3.2; 6) takes an even more restrictive view. He points out (correctly) that an utterance such as (28) may be used to reject not only previous utterances but also attributed thoughts or assumptions which are "in the discourse model"; however, instead of concluding that all these cases are echoic denials, he decides to exclude all of them from his category of echoic use. For him, the only genuine cases of echoic denial are metalinguistic, based on resemblances in form. Carston (1996) offers what seems to me a more satisfactory account. She includes in the category of echoic denials the full set of cases involving both attributed utterances and attributed thoughts. On her account, (28)-(31) would all be treated as echoic, as would any utterance used to metarepresent and reject an attributed utterance or thought. (For discussion, see Burton-Roberts 1989; Carston 1996; Horn 1989; Iwata 1998; McCawley 1991; Noh 1998a; van der Sandt 1991).

"Echo questions" are formally distinguishable from regular interrogatives by their declarative syntax and rising intonation. Their treatment has generally run parallel to the treatment of metalinguistic negation. Consider (22b) above, or (32b)-(34b):

(32) a. *Peter*: You finally managed to solve the problems.
b. *Mary*: Managed? I solved them in two minutes. (Noh 1998a, 218)

(33) a. *Peter*: I need a holiday.
b. *Mary*: You need a holiday? What about me?

(34) a. *Tourist*: Where can I find some tom[eiDouz]?
b. *Londoner*: You want tom[eiDouz]? Try New York.

All four questions are clearly echoic in the sense defined above: the speaker echoes and questions some aspect of the form or content of an attributed utterance. However, as with echoic denials, linguistic analyses of echo questions have generally been over-restrictive. "Echo questions" are generally defined as echoing prior utterances, not thoughts:

> Echo questions are distinguished from other questions by their restricted context. An echo occurs in dialogue as a reaction to a prior utterance and is interpretable only with respect to it, while other questions may be the first or the only utterance in a discourse. (Banfield 1982, 124)

> Echo questions generally require a linguistic context in which the original utterance ... has been previously uttered within the discourse. (Horn 1989, 381)

Yet there seem to be clear cases of echoic questions used to metarepresent attributed thoughts. Compare (35a-c):

(35) a. *Mary* [seeing Peter walk towards the door]: Just a minute. You're going shopping?
b ?*Mary* [seeing Peter walk towards the door]: Just a minute. Henry VIII had six wives?
c. *Mary* [seeing Peter walk towards the door]: Just a minute. Did Henry VIII have six wives?

Here, the echoic question in (35a) and the regular interrogative in (35c) are pragmatically appropriate, but the echoic question in (35b) is not. The obvious way of explaining this would be to treat the utterances in (35a) and (35b) as echoing and questioning thoughts that Mary attributes to Peter. In (35a), his behaviour gives her ground for inferring his thoughts even though he hasn't spoken: in (35b), it does not.

This would enable us to maintain the parallel between verbal irony, denials and echo questions: all are tacitly attributive, and all may be used for the attribution both of utterances and of thoughts (Blakemore 1994; Escandell-Vidal 1998; Noh 1998a, 1999).

4.4 Non-attributive cases

So far, the only lower-order representations I have looked at have been attributive. As noted in section 1, there are also non-attributive cases: mentions of sentence types, utterance types or proposition types, as in (6) (repeated below):

(6) a. 'Dragonflies are beautiful' is a sentence of English.
b. 'Shut up' is rude.
c. It's true that tulips are flowers.
d. *Roses and daisies are flowers* entails that roses are flowers.
e. I like the name 'Petronella'.
f. 'Abeille' is not a word of English.
g. *Tulip* implies *flower*.

These are worth considering because they are not obviously linked to the metapsychological or metacommunicative abilities, and might contrast in interesting ways with attributions of utterances and thoughts.

To understand the cases of mention in (6), the hearer must be able to recognise linguistic, logical or conceptual resemblances between representations considered in the abstract rather than tied to a particular individual, place or time. Because no attribution is involved, this ability might be present even if the intuitive metapsychological or metacommunicative capacity is impaired. Indeed, there is evidence from the autobiographical writings of people with autism (several of whom have been students of linguistics) of a serious interest in linguistic form (Frith & Happé 1999; Willey 1999; Williams 1992). Here is how one of them describes her fascination with language:

> Linguistics and the act of speaking itself have always been among my keenest interests … Words, and everything about them, hold my concentration like nothing else. On my overstuffed bookshelf sit several thesauruses, a half dozen dictionaries, famous quotations books, and a handful of personal reflection journals. Language appeals to me because it lends itself to rules and precision even more often than it does to subjectivity … Some words can please my eyes, given that they have the symmetry of line and shape I favor. Other words can fascinate me by the melodies they sing when they are spoken. Properly handled … words can work miracles on

> my sensibilities and my understanding of the world, because each one has its own personality and nuance and its own lesson to teach. (Willey 1999, 30)

This is in marked contrast to the comments of the same writer on her inability to discern the intentions behind other people's use of words.

Relevance theorists have argued that there are several further types of non-attributive metarepresentation which have been less widely recognised, and which clearly contrast with the attributive cases discussed above. For example, regular (non-attributive) interrogatives and exclamatives have been treated in relevance theory as representations of *desirable thoughts* (or desirable information); and regular (non-attributive) negations and disjunctions have been treated as representations of *possible thoughts* (possible information). Parallel to these, we might expect to find (non-attributive) representations of possible and desirable utterances. I will end this survey with a few illustrations of each.

In relevance theory, regular (non-echoic) interrogatives such as those in (36) have been treated as the metarepresentational counterparts of imperatives:

(36) a. Is today Tuesday?
b. What day is it today?
c. When are we leaving?

Imperatives represent *desirable states of affairs*; interrogatives represent *desirable thoughts*. Someone who utters an imperative is thinking about a state of affairs, which she regards as desirable from someone's point of view. Someone who utters an interrogative is thinking about a thought (or item of information), which she regards as desirable from someone's point of view. Since information can be desirable only because it is relevant, this amounts to claiming that interrogatives represent relevant answers (Clark 1991; Sperber & Wilson 1986/1995, chapter 4, section 10; Wilson & Sperber 1988a).

This account of interrogatives has some advantages over alternative analyses. For example, in speech-act theory, interrogatives are generally treated as encoding requests for information (Bach and Harnish 1979; Harnish 1994; Searle 1969). One problem with the speech-act approach is that not all interrogative utterances are requests for information: they may be offers of information, rhetorical questions, exam questions, idle speculations, and so on. The relevance-theoretic account solves this problem in the following way. An interrogative utterance merely indicates that the speaker regards the answer as relevant to *someone*. It will only be understood as a request for information if two further contextual conditions are fulfilled: (a) the speaker regards the answer as

relevant to herself; and (b) the hearer is in a position to provide it. In other conditions, it will be differently understood. For example, it will be understood as an offer of information if (a) the speaker regards the answer as relevant to the hearer, and (b) the speaker herself is in a position to provide it. Other types of interrogative speech act also fall out naturally from this account (Clark 1991; Sperber & Wilson 1986/1995, chapter 4, section 10; Wilson & Sperber 1988a, 1988b).

The analysis of interrogatives as inherently metarepresentational brings them into interesting contact with the literature on mindreading. On the relevance-theoretic account, the production and interpretation of interrogatives necessarily involves a higher order of metarepresentational ability than standard declaratives, but differs in two respects from examples that have been central in the "theory-of-mind" literature: first, the metarepresented proposition is not attributive, and second, it is not treated as either false (as in the false-belief task) or true (as in pretence). The fact that it is not attributive means that we might expect someone who fails second-order "theory of mind" tasks to pass tasks involving regular interrogatives, but fail on echo questions, for example.

If interrogatives metarepresent desirable thoughts, we might expect to find utterances used to metarepresent desirable utterances. There is no shortage of candidates. Here are some possible examples:

(37) a. *Vicar to bride*: I, Amanda, take you, Bertrand, to be my lawful wedded husband.
b. *Bride*: I, Amanda, take you, Bertrand, to be my lawful, wedded husband.

(38) *Mary to Peter* (as doorbell rings): If that's John, I'm not here. (Noh 1998a)

(39) a. *Quiz-show host*: The first man to walk on the moon was?
b. *Contestant*: Neil Diamond.

In (37a), the vicar metarepresents an utterance that he wants the bride to produce. In the "speech-act conditional" in (38), the consequent is used to metarepresent an utterance that Mary wants Peter to produce; (38) expresses something like the proposition "If that's John, *say* I'm not here" (Sweetser 1990; Noh 1998b; van der Auwera 1986). The quiz-show question in (39) might be analysed on similar lines: the host is not producing a regular interrogative, but metarepresenting an utterance he wants the contestant to produce. Further illustrations include the utterances of prompters in a theatre, and solicitors whispering answers to their clients in court.

The literature on standard mentions contains many examples of utterances used to represent possible thoughts and utterances. For instance, mentions of propositions, as in (6c-d) above, amount to metarepresentations of possible thought types, and mentions of

utterances, as in (6b) above, amount to metarepresentations of possible utterance types. But there may be a much wider range of candidates. Consider the examples in (40):

(40) a. Ducks don't bite.
b. Maybe I'll leave.
c. Either William will become a soldier or Harry will.

Regular (non-attributive) uses of negation, modals and disjunctions, as in (40a-c), seem to presuppose the ability to think about possible thoughts and evaluate their truth or falsity, which would make all these examples metarepresentational. The development of attributive and non-attributive negatives, modals and interrogatives contrast in interesting ways, shedding light on the interaction between metalogical, metacommunicative and metapsychological abilities (Bloom 1991; Gombert 1990; Morris & Sloutsky 1998; Noveck, Ho & Sera 1996; Overton 1990; Papafragou 1998c).

Apart from standard mentions of utterance types, as in (6b) above, metarepresentations of possible utterances might include the following advertisements from a recent Glenfiddich whisky campaign in England:

(41) *Picture of a newspaper with the headline*: French admit Britain is best.
Caption: Till then, there's Glenfiddich to enjoy.

(42) *Picture of a newspaper with the headline*: World's funniest man is a Belgian.
Caption: Till then, there's Glenfiddich to enjoy.

Drafts, essay plans and rehearsals of future conversations might provide further examples. There is evidence, then, that all four categories of non-attributive representation are filled.

5 Conclusion

This survey was designed to show something of the depth and variety of the metarepresentational abilities used in verbal comprehension. Language is full of metarepresentational devices, which are often quite fragmentary or incomplete: I have argued that they provide no more than triggers for spontaneous metacommunicative processes by which speaker meanings are inferred. I have outlined a pragmatic comprehension procedure which might help to resolve indeterminacies in meaning and form the basis for a modularised metacommunicative ability, itself a sub-part of a more general metapsychological ability, or "theory of mind".

The processing of linguistic metarepresentations also interacts in more specific ways with the metapsychological ability. As I have shown, linguistic metarepresentations vary both in degree of explicitness and in the type of original they are used to represent: utterances, thoughts or abstract representations. By comparing comprehension in these different types of case, it might be possible to gain new insight into the metapsychological and metacommunicative abilities. To take just one example, there are cases (such as (20) above) in which the mindreading ability directly feeds the comprehension process, by interpreting paralinguistic information (gestures, facial expressions, intonation, and so on) to provide information about the speaker's mood or epistemic state, which may in turn be picked out by the pragmatic comprehension procedure and attributed as part of a speaker's meaning. Inferring these aspects of speaker's meaning is likely to prove particularly difficult for people whose general mindreading ability is weak. It would be interesting to check whether the use of overt linguistic devices would facilitate comprehension, and if so, in what way.[4] For example, is it easier to attribute a false belief when expressed or implied by an utterance ("The ball is in the cupboard") than by inferring it on the basis of non-communicative behaviour?

From the linguist's point of view, there are also benefits to be gained by considering metarepresentational devices in the context of the more general metapsychological and metacommunicative abilities. As I have shown, in studying linguistic metarepresentations, linguists have tended to concentrate on cases involving the attribution of utterances, and many echoic utterances and "hearsay" devices may have broader uses than existing research suggests. Studies of lexical acquisition are already being conducted within a broader metacommunicative and metapsychological framework; the acquisition of specifically metarepresentational devices within this framework should also yield interesting results.

References

Astington, J., P. Harris & D. Olson (eds.). (1988). *Developing Theories of Mind.* Cambridge: Cambridge University Press.

Bach, K. & R. M. Harnish (1979). *Linguistic Communication and Speech Acts.* Cambridge, MA: MIT Press.

Banfield, A. (1982). *Unspeakable Sentences: Narration and Representation in the Language of Fiction*. Boston: Routledge & Kegan Paul.

Baron-Cohen, S. (1995). *Mindblindness: An Essay on Autism and Theory of Mind.* Cambridge, MA: MIT Press.

[4] Even overtly metarepresentational devices may leave a lot to be inferred, as the example with "disappointed" shows.

Baron-Cohen, S., A. Leslie & U. Frith (1985). Does the autistic child have a 'theory of mind'? *Cognition* 21: 37-46.

Baron-Cohen, S., H. Tager-Flusberg & D. Cohen (eds.). (1993). *Understanding Other Minds: Perspectives from Autism*. Oxford: Oxford University Press.

Bartsch, K. & H. Wellman (1995). *Children Talk about the Mind.* Oxford: Oxford University Press.

Bezuidenhout, A. & M. S. Sroda (1996). Disambiguation of reference by young children: An application of relevance theory. Paper presented at Relevance Theory Workshop, Tokyo, March 1996.

Bezuidenhout, A. & M. S. Sroda (1998). Children's use of contextual cues to resolve referential ambiguity: An application of relevance theory. *Pragmatics and Cognition* 6.

Blakemore, D. (1991). Performatives and parentheticals. *Proceedings of the Aristotelian Society* 91: 197-214.

Blakemore, D. (1994). Echo questions: A pragmatic account. *Lingua* 4:197-211.

Blakemore, D. (1997). Restatement and exemplification: A relevance-theoretic reassessment of elaboration. *Pragmatics and Cognition* 5: 1-19.

Blass, R. (1989). Grammaticalisation of interpretive use: the case of re in Sissala. *Lingua* 79: 229-326.

Blass, R. (1990). *Relevance Relations in Discourse: A Study with Special Reference to Sissala.* Cambridge: Cambridge University Press.

Bloom, L. (1991). *Language Development from Two to Three.* Cambridge: Cambridge University Press.

Bloom, P. (1997). Intentionality and word learning. *Trends in Cognitive Sciences* 1: 9-12.

Booth, W. (1961). *The Rhetoric of Fiction*. Chicago, IL: Chicago University Press.

Bretherton, I. (1991). Intentional communication and the development of an understanding of mind. In Frye & Moore (eds.). 49-75.

Bretherton, I. & M. Beeghly (1982). Talking about internal states: The acquisition of an explicit theory of mind. *Developmental Psychology* 18: 906-21.

Burton-Roberts, N. (1989). *The Limits to Debate: A Revised Theory of Semantic Presupposition.* Cambridge: Cambridge University Press.

Cappelen, H. & E. Lepore (1997a). Varieties of quotation. *Mind* 106: 429-450.

Cappelen, H. & E. Lepore (1997b). On an alleged connection between indirect speech and theory of meaning. *Mind and Language* 12: 278-96.

Carston, R. (1988). Implicature, explicature and truth-theoretic semantics. In R. Kempson (ed.). 155-81. Reprinted in S. Davis (ed.). 33-51.

Carston, R. (1996). Metalinguistic negation and echoic use. *Journal of Pragmatics* 25: 309-330.

Carston, R. (1998). *Pragmatics and the Explicit-Implicit Distinction*. University College London PhD thesis. To be published by Blackwell, Oxford.

Chafe, W. & J. Nichols (eds.). (1986). *Evidentiality: The Linguistic Coding of Epistemology*. Norwood, N.J.: Ablex.

Clark, B. (1991). *Relevance Theory and the Semantics of Non-declaratives*. University of London PhD thesis.

Clark, H. & T. Carlson (1981). Context for comprehension. In J. Long & A. Baddeley (eds.). *Attention and performance* IX. 313-30. Hillsdale, NJ: Lawrence Erlbaum.

Clark, H. & R. Gerrig (1984). On the pretense theory of irony. *Journal of Experimental Psychology: General* 113: 121-6.

Clark, H. & R. Gerrig (1990). Quotations as demonstrations. *Language* 66: 764-805.

Clark, H. & C. Marshall (1981). Definite reference and mutual knowledge. In Aravind J., B. Webber & I. Sag (eds.). *Elements of Discourse Understanding*. 10-63. Cambridge: Cambridge University Press.

Cohn, D. (1978). *Transparent Minds: Narrative Modes for Presenting Consciousness in Fiction*. Princeton: Princeton University Press.
Coulmas, F. (ed.). (1986). *Direct and Indirect Speech*. Berlin: Mouton de Gruyter.
Curcó, C. (1998). Indirect echoes and verbal humour. In Rouchota & Jucker (eds.). 305-325.
Davies, M. & T. Stone (eds.). (1995a). *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.
Davies, M. & T. Stone (eds.). (1995b). *Folk Psychology*. Oxford: Blackwell.
Davidson, D. (1968). On saying that. In D. Davidson (1984). 93-108.
Davidson, D. (1979). Quotation. In D. Davidson (1984). 79-92.
Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.
Davis, S. (ed.). (1991). *Pragmatics: A Reader*. Oxford: Oxford University Press.
Ducrot, O. (1983). Puisque: essai de description polyphonique. In M. Herslund et al. (eds.). *Analyse grammaticale du français*. Special issue of *Revue romane* 24: 166-85.
Escandell-Vidal, V. (1998). Metapropositions as metarepresentations. Paper delivered to the Relevance Theory Workshop, Luton, September 1998.
Fillmore, Ch. (1990). Epistemic stance and grammatical form in English conditional sentences. *Chicago Linguistic Society* 26: 137-62.
Fodor, J. (1992). A theory of the child's theory of mind. *Cognition* 44: 283-96.
Fludernik, M. (1993). *The Fictions of Language and the Languages of Fiction: The Linguistic Representation of Speech and Consciousness*. London: Routledge.
Frith, U. & F. Happé (1999). Theory of mind and self-consciousness: What is it like to be autistic? *Mind and Language* 14: 1-22.
Frye, D. & Ch. Moore (eds.). (1991). *Children's Theories of Mind: Mental States and Social Understanding*, 49-75. Hillsdale, NJ: Lawrence Erlbaum.
Garnham, A. & J. Perner 1990. Does manifestness solve problems of mutuality? *Behavioral and Brain Sciences* 13: 178-9.
Garver, N. (1965). Varieties of use and mention. *Philosophy and Phenomenological Research* 26: 230-38.
Gibbs, R. (1987). Mutual knowledge and the psychology of conversational inference. *Journal of Pragmatics* 11: 561-88.
Gibbs, R. (1994). *The Poetics of Mind: Figurative Thought, Language and Understanding*. New York: Cambridge University Press.
Gombert, J. (1990). *Metalinguistic Development*. Hassocks, Sussex: Harvester Press.
Gopnik, A. & H. Wellman (1992). Why the child's theory of mind really is a theory. *Mind and Language* 7: 145-71.
Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (eds.). *Syntax and Semantics 3: Speech Acts*. 41-58. Reprinted in Grice (1989). 22-40.
Grice, H. P. (1982). Meaning revisited. In N. Smith (ed.). 223-43.
Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
Gutt, E.-A. (1991). *Translation and Relevance: Cognition and Context*. Oxford: Blackwell.
Happé, F. (1993). Communicative competence and theory of mind in autism: A test of relevance theory. *Cognition* 48: 101-19.
Happé, F. (1994). *Autism: An Introduction to Psychological Theory*. Cambridge MA: Harvard University Press.
Harnish, R. M. (1994). Mood, meaning and speech acts. In S. Tsohatzidis (ed.). *Foundations of Speech-act Theory*. London: Routledge.
Horn, L. (1985). Metalinguistic negation and pragmatic ambiguity. *Language* 61: 121-174.

Horn, L. (1989). *A Natural History of Negation*. Chicago, IL: Chicago University Press.
Ifantidou-Trouki, E. (1993). Sentence adverbials and relevance. *Lingua* 90: 69-90.
Ifantidou, E. (1994). *Evidentials and Relevance*. University College London, PhD thesis.
Itani, R. (1996). *Semantics and Pragmatics of Hedges in English and Japanese*. Tokyo: Hituzi Syobo.
Iwata, S. (1998). Some extensions of the echoic analysis of metalinguistic negation. *Lingua* 105: 49-65.
Jorgensen, J., G. Miller & D. Sperber (1984). Test of the mention theory of irony. *Journal of Experimental Psychology: General* 113: 112-20.
Kasher, A. (ed.). (1998). *Pragmatics: Critical concepts*, vols. I-VI. London: Routledge.
Kempson, Ruth (ed.). (1988). *Mental Representations: The Interface between Language and Reality*. Cambridge: Cambridge University Press.
Kiparsky, C. & P. Kiparsky (1971). Fact. In D. Steinberg & L. Jakobovits (eds.). *Semantics: An Interdisciplinary Reader*. 345-69. Cambridge: Cambridge University Press.
Kreuz, R. & S. Glucksberg (1989). How to be sarcastic: The echoic reminder theory of irony. *Journal of Experimental Psychology: General* 118: 374-86.
Kumon-Nakamura, S., Glucksberg, S. & M. Brown (1995). How about another piece of pie: the allusional pretense theory of discourse irony. *Journal of Experimental Psychology: General* 124: 3-21.
Lascelles, M. (1939/1965). *Jane Austen and her Art*. Oxford: Oxford University Press.
Leech, G. & M. Short (1981). *Style in Fiction: An Introduction to English Fictional Prose*. London: Longman.
Leslie, A. (1987). Pretense and representation: The origins of "theory of mind". *Psychological Review* 94: 412-426.
Leslie, A. (1991). The theory of mind impairment in autism: Evidence for a modular mechanism of development? In A. Whiten (ed.). 63-78.
Leslie, A. & F. Happé (1989). Autism and ostensive communication: The relevance of metarepresentation. *Development and Psychopathology* 1: 205-12.
Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
Lewis, C. & P. Mitchell (eds.). (1994). *Children's Early Understanding of Mind: Origins and Development*. Hillsdale, NJ: Lawrence Erlbaum.
Lock, A. (1980). *The Guided Reinvention of Language*. New York, NY: Academic Press.
Martin, R. (1992). Irony and universe of belief. *Lingua* 87: 77-90.
McCawley, J. (1991). Contrastive negation and metalinguistic negation. *Chicago Linguistic Society 27: Parasession on Negation*. 189-206.
McHale, B. (1978). Free indirect discourse: a survey of recent accounts. *PTL: Journal for Descriptive Poetics and Theory of Literature* 3: 235-87.
Morris, A. & V. Sloutsky (1998). Understanding of logical necessity: Developmental antecedents and cognitive consequences. *Child Development* 69: 721-741.
Neale, S. (1992). Paul Grice and the philosophy of language. *Linguistics and Philosophy* 15: 509-59.
Noh, E.-J. (1998a). *The Semantics and Pragmatics of Metarepresentations in English: A Relevance-Theoretic Approach*. University College London, PhD thesis. To be published by John Benjamins, Amsterdam.
Noh, E.-J. (1998b). A relevance-theoretic account of metarepresentative uses in conditionals. In V. Rouchota & A. Jucker (eds.). 271-304.
Noh, E.-J. (1999). Echo questions: Metarepresentation and pragmatic enrichment. *Linguistics and Philosophy*.

Noveck, I., S. Ho & M. Sera (1996). Children's understanding of epistemic modals. *Journal of Child Language* 23: 621-643.
Nuti, M. (in preparation). *Pragmatics and mindreading*. University College London PhD thesis.
Origgi, G. & D. Sperber (forthcoming). Issues in the evolution of human language and communication. To appear in P. Carruthers & A. Chamberlain (eds.). *Evolution and the Human Mind: Language, Modularity and Social Cognition*. Cambridge: Cambridge University Press.
Overton, W. (ed.). (1990). *Reasoning, necessity and logic: Developmental perspectives*. Hillsdale, NJ: Lawrence Erlbaum.
Papafragou, A. (1998a). *Modality and the Semantics-Pragmatics Interface*. University College London PhD thesis. To be published by Elsevier.
Papafragou, A. (1998b). Inference and word meaning: The case of modal auxiliaries. *Lingua* 105: 1-47.
Papafragou, A. (1998c). The acquisition of modality: Implications for theories of semantic representation. *Mind and Language* 13: 370-399.
Partee, B. (1973). The syntax and semantics of quotations. In S. Anderson & P. Kiparsky (eds.). *A Festschrift for Morris Halle*. 410-418. New York, NY: Holt, Rinehart and Winston.
Recanati, F. (1986). On defining communicative intentions. *Mind and Language* 1: 213-42.
Recanati, F. (1987). *Meaning and Force.* Cambridge: Cambridge University Press.
Recanati, F. (1989). The pragmatics of what is said. *Mind and Language* 4: 295-329. Reprinted in S. Davis (ed.). 97-120.
Rouchota, V. & A. Jucker (eds.). (1998). *Current Issues in Relevance Theory.* Amsterdam: John Benjamins.
Saka, P. (1998). Quotation and the use-mention distinction. *Mind* 107: 113-35.
Schiffer, S. (1972). *Meaning*. Oxford: Clarendon Press.
Scholl, B. & A. Leslie (1999). Modularity, development and "theory of mind". *Mind and Language* 14: 131-153.
Searle, J. (1969). *Speech Acts*. Cambridge: Cambridge University Press.
Smith, N. (ed.). (1982). *Mutual Knowledge.* London: Academic Press.
Smith, N. & I.-M. Tsimpli (1995). *The Mind of a Savant: Language Learning and Modularity*. Oxford: Blackwell.
Sperber, D. (1984). Verbal irony: Pretense or echoic mention? *Journal of Experimental Psychology: General* 113: 130-6.
Sperber, D. (1994). Understanding verbal understanding. In J. Khalfa (ed.). *What is Intelligence*? 179-98. Cambridge: Cambridge University Press.
Sperber, D. (1996). *Explaining Culture: A Naturalistic Approach*. Oxford: Blackwell.
Sperber, D. (1997). Intuitive and reflective beliefs. *Mind and Language* 12: 67-83.
Sperber, D. (forthcoming). Metarepresentations in an evolutionary perspective. To appear in S. Davis & D. Sperber (eds.). *Proceedings of the Tenth Vancouver Cognitive Science Conference.*
Sperber, D. & D. Wilson (1981). Irony and the use-mention distinction. In P. Cole (ed.). *Radical Pragmatics*. 295-318. New York, NY: Academic Press.
Sperber, D. & D. Wilson (1986/1995). *Relevance: Communication and Cognition*. Oxford: Blackwell.
Sperber, D. & D. Wilson (1987). Précis of *Relevance: Communication and Cognition. Behavioral and Brain Sciences* 10: 697-754. Reprinted in A. Kasher (ed.). vol. V, 82-115.
Sperber, D. & D. Wilson (1990a). Spontaneous deduction and mutual knowledge. *Behavioral and Brain Sciences* 13: 179-84.
Sperber, D. & D. Wilson (1990b). Rhetoric and relevance. In J. Bender & D. Wellbery (eds.). *The Ends of Rhetoric: History, Theory, Practice.* 140-56. Stanford, CA: Stanford University Press.

Sperber, D. & D. Wilson (1998a). Irony and relevance A reply to Seto, Hamamoto and Yamanashi. In R. Carston & S. Uchida (eds.). *Relevance Theory: Applications and Implications*. 283-93. Amsterdam: John Benjamins.

Sperber, D. & D. Wilson (1998b). The mapping between the mental and the public lexicon. In P. Carruthers & J. Boucher (eds.). *Language and Thought*: *Interdisciplinary Themes*. 184-200. Cambridge: Cambridge University Press.

Sternberg, M. (1982a). Proteus in quotation-land: Mimesis and the forms of reported discourse. *Poetics Today* 3: 107-56.

Sternberg, M. (1982b). Point of view and the indirections of direct speech. *Language and Style* 15: 67-117.

Sweetser, E. (1990). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge: Cambridge University Press.

Tomasello, M. & A. C. Kruger. (1992). Acquiring verbs in ostensive and non-ostensive contexts. *Journal of Child Language* 19: 311-333.

Urmson, J. (1963). Parenthetical verbs. In C. Caton (ed.). *Philosophy and Ordinary Language*. 220-40. Urbana: University of Illinois Press.

van der Auwera, J. (1986). Conditionals and speech acts. In E. Traugott, A. ter Meulen, J. Reilly & Ch. Ferguson (eds.). *On Conditionals*. 197-214. Cambridge: Cambridge University Press.

van der Sandt, R. (1991). Denial. *Chicago Linguistic Society 27: Parasession on Negation*. 331-44.

Wade, E. & H. Clark (1993). Reproduction and demonstration in quotations. *Journal of Memory and Language* 32: 805-19.

Walton, K. (1976). Points of view in narrative and depictive representation. *Nous* 10: 49-61.

Walton, K. (1990). *Mimesis as Make-believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press.

Wellman, H. (1990) *The Child's Theory of Mind.* Cambridge MA: MIT Press.

Whiten, A. (ed.). (1991). *Natural Theories of Mind: Evolution, Development and Simulation of Everyday Mindreading*. Oxford: Blackwell.

Willey, L. H. (1999). *Pretending to be Normal: Living with Asperger's Syndrome*. London: Jessica Kingsley Publishers.

Williams, D. (1992). *Nobody Nowhere*. New York, NY: Avon Books.

Wilson, D. & D. Sperber (1988a). Mood and the analysis of non-declarative sentences. In J. Dancy, J. Moravcsik & Ch. Taylor (eds.). *Human Agency: Language, Duty and Value*. 77-101. Stanford, CA: Stanford University Press. Reprinted in A. Kasher (ed.). vol. II, 262-89.

Wilson, D. & D. Sperber (1988b). Representation and relevance. In R. Kempson (ed.). *Mental Representation*: *The Interface between Language and Reality*. 133-53. Cambridge: Cambridge University Press.

Wilson, D. & D. Sperber (1992). On verbal irony. *Lingua* 87: 53-76.

Wilson, D. & D. Sperber (1993). Linguistic form and relevance. *Lingua* 90: 1-25.

Wilson, D. & D. Sperber (in preparation). Truthfulness and relevance. Paper delivered to the San Marino conference on Grice's Heritage, May 1998.

Winner, E. (1988). *The Point of Words: Children's Understanding of Metaphor and Irony*. Cambridge MA: Harvard University Press.