Relevance theory[1]

Nicholas Allott

ABSTRACT

Relevance theory is a framework for the study of cognition, proposed primarily in order to provide a psychologically realistic account of communication. This paper i) presents relevance theory's central commitments in detail and explains the theoretical motivations behind them; and ii) shows some of the ways in which these core principles are brought to bear on empirical problems.

The core of relevance theory can be divided into two sets of assumptions. Assumptions relating to cognition in general include the definition of relevance as a trade-off between effort and effects, and the claim that cognition tends to maximise relevance. Assumptions about communication include the claims that understanding an utterance is a matter of inferring the speaker's communicative and informative intentions; and that the communicative principle of relevance and the presumption of optimal relevance mandate the relevance-theoretic comprehension procedure, a heuristic that guides the search for the intended interpretation of utterances. Relevance theorists model communication in terms of the working of this comprehension procedure. There are, in addition, several strategies that guide the explanation of phenomena in relevance theory, including: i) a stronger form of Grice's Modified Occam's Razor, ii) the possibility of dividing what is linguistically encoded between conceptual and procedural information; iii) the interpretive/descriptive distinction; iv) the use of *ad hoc* concepts.

1 Introduction: the relevance-theoretic research programme

Relevance theory is a rather wide-ranging framework (or 'research programme' – see below) for the study of cognition, devised primarily in order to provide an account of communication that is psychologically realistic and empirically plausible. It was originally proposed by Sperber and Wilson (1986b; 1987). Other key publications include Blakemore, 1987; Sperber & Wilson, 1995; Carston, 2002 and Wilson & Sperber, 2012.

For some time relevance theory has been one of the leading programmes of research in   pragmatics. There has been work within the relevance-theoretic framework[2] on such central topics as scalar implicatures (Carston, 1998; Breheny, Katsos & Williams, 2006; Noveck & Sperber, 2007), bridging (Wilson & Matsui, 1998; Matsui, 2000), speech acts and mood (Sperber & Wilson, 1986b, pp. 243–254; Wilson & Sperber, 1988; Jary, 2007; Jary, 2010), disambiguation (Sperber & Wilson, 1986b, pp. 183–193), discourse particles (Blakemore, 1987; Blakemore, 2000; Blakemore, 2002; Blakemore, 2004; Iten, 2005),

---

2 The references given here are far from exhaustive. For many more references, sorted by author and by subject matter, see Francisco Yus' online relevance theory bibliography at http://www.ua.es/personal/francisco.yus/rt.html

evidentials (Ifantidou, 2001), loose talk (Sperber & Wilson, 1986a; Carston, 1997a; Wilson & Sperber, 2002), literary language (Sperber & Wilson, 1986b, ch. 4; Clark, 1996; Pilkington, 2000; Sperber & Wilson, 2008), genre (Unger, 2006), translation (Gutt, 1991), non-verbal communication (Wharton, 2009), the referential/attributive distinction (Rouchota, 1992; Bezuidenhout, 1997; Powell, 2001; Powell, 2010) and rhetorical tropes such as metaphor (Sperber & Wilson, 1986b, pp. 231–237; Carston, 1997a; Vega, 2007; Sperber & Wilson, 2008; Carston, 2010b) and irony (Sperber & Wilson, 1981; Sperber & Wilson, 1986b, pp. 237–243; Wilson & Sperber, 1992; Sperber & Wilson, 1998a; Wilson, 2006). The theory has had considerable influence in the disputed borderlands between semantics, pragmatics and philosophy of language, including ongoing debates about the distinction between what is explicitly and what implicitly communicated, and the extent to which pragmatic inference affects the proposition expressed by an utterance (Wilson & Sperber, 1981; Sperber & Wilson, 1986b, ch 4; Carston, 1988; Carston, 2002, ch. 2–5; Carston, 2010a). Relevance theory has also inspired considerable work on the application of experimental and developmental evidence to pragmatics and related questions in the psychology of reasoning, helping to shape the emerging field of experimental pragmatics (Jorgensen, Miller & Sperber, 1984; Happé, 1993; Sperber, Cara & Girotto, 1995; Bezuidenhout & Sroda, 1998; Nicolle & Clark, 1999; van der Henst, Carles & Sperber, 2002; van der Henst, Politzer & Sperber, 2002; Happé & Loth, 2002; Noveck & Sperber, 2004; Breheny, Katsos & Williams, 2006; Noveck & Sperber, 2007; Chevallier, Wilson, Happé & Noveck, 2010; Chevallier, Noveck, Happé & Wilson, 2011).

Despite its reach and popularity, however, relevance theory is poorly understood beyond its practitioners. There is confusion among both linguists and philosophers about what relevance theorists are committed to and what kinds of explanations they attempt to give[3].

This paper attempts to clarify these issues by i) presenting relevance theory's central commitments in detail and explaining the theoretical motivations behind them; and ii) showing some of the ways in which these core principles are brought to bear on empirical problems.

As Wilson and Sperber say:

> Like other psychological theories, [relevance theory] has testable consequences: it can suggest experimental research, and is open to confirmation, disconfirmation, or fine-tuning in the light of experimental evidence. As with other theories of comparable scope, its most general claims can only be tested indirectly. For example, the Cognitive Principle of Relevance suggests testable predictions only when combined with descriptions of particular cognitive mechanisms (e.g. for perception, categorization. memory, or inference). (Wilson & Sperber, 2004, pp. 625–626)

There are echoes here of the model of scientific research proposed by Imre Lakatos. According to Lakatos, scientists work within competing research programmes, and each research programme has two components: i) a 'hard core' of fundamental theoretical commitments; and ii) auxiliary hypotheses (Lakatos, 1968, p. 168 ff.)[4]. Most of the predictions relevance theory makes do not follow from the hard core on its own: it is only once the auxiliary hypotheses are also taken into account that the theory (or, rather, the programme as a whole) makes predictions[5]. Each research programme also has a positive heuristic, specifiying strategies for forming theories outside of the hard core: i.e. suggesting what 'paths

---

3 See Sperber and Wilson's replies to comments on their precis of 'Relevance' in Behavioral and Brain Science (Sperber & Wilson, 1987), and Wedgewood (2007) and Kjøll (2010) who have argued that certain recent criticisms of relevance theory in the philosophy of language literature are based on fundamental misunderstandings about relevance theory's commitments.

4 See also Lakatos, 1970. Lakatos' papers on the methodology of science are collected in volume 1 of Lakatos, Worrall & Currie, 1978. For critical commentary see Hacking, 1979.

5 Lakatos also claims that the core commitments are to be kept, while auxiliary hypotheses should be modified or disposed of in response to empirical challenges. (He calls this the 'negative heuristic': Lakatos, 1968, p. 169.) I return to this point briefly in the conclusion of this paper, where I discuss some changes that have occurred in the core of relevance theory.

of research' to pursue (Lakatos, 1968, p. 168).

This paper describes the central assumptions of relevance theory in detail and then sketches some of the strategies that relevance theorists use in developing theories beyond that core.

1.1 The central assumptions and positive heuristic of relevance theory

The core of relevance theory can be divided into two sets of assumptions. Assumptions in the first set relate to cognition in general, assumptions in the second to communication more specifically, particularly to utterance interpretation.

The central assumptions that relevance theory makes about human cognition include the definition of relevance as a trade-off between effort and effects; the cognitive principle of relevance, which is the claim that cognition tends to maximise relevance; and the views, shared with other work in cognitive science, that cognition is a matter of (or at least can be well modelled as) computations over mental representations, and that human beings possess a 'deductive device' which plays a central role in spontaneous inference. I set out these core assumptions relating to cognition in section 2 of this paper.

The core of relevance theory as it relates specifically to communication includes the Gricean claim that understanding an utterance is a matter of inferring what the speaker intended to convey from what she utters (in what way, in what circumstances). Another fundamental of relevance theory, departing somewhat from Grice, is that there are exactly two speaker's intentions that are central to communication, namely the informative intention and the communicative intention. The last main part of the hard core relating specifically to communication is entirely original to relevance theory: the communicative principle of relevance and the presumption of optimal relevance, which mandate the relevance-theoretic comprehension procedure, a heuristic that guides the search for the correct (i.e. intended) interpretation of utterances. I examine the core assumptions that are specific to communication in section 3.

The characteristic approach of relevance theory to the explanation of communicative phenomena is a corollary of its central commitments. Relevance theorists try to give psychologically realistic explanations and to understand communicated meaning in terms of the working of the relevance-theoretic comprehension procedure. This way of working is at the heart of relevance theory's 'positive heuristic', but in section 4 I show that there are several additional strategies that guide the explanation of phenomena in relevance theory including: i) Grice's Modified Occam's Razor, in a stronger form; ii) the possibility of dividing what is linguistically encoded between conceptual and procedural information; iii) the interpretive/descriptive distinction; iv) the use of *ad hoc* concepts.

2 Relevance theory and cognition

The central claim of relevance theory is that, as a result of constant selection pressures, the human cognitive system has developed a variety of dedicated (innate or acquired) mental mechanisms or biases which tend to allocate attention to inputs with the greatest expected relevance, and process them in the most relevance-enhancing way. (Wilson, 2009, p. 394)

2.1 The cognitive principle of relevance

At the centre of the hard core of relevance theory are the cognitive principle of relevance and the definition of relevance as a trade-off of cognitive benefit against processing cost. The cognitive principle is the hypothesis that cognitive systems tend to maximise relevance.

Cognitive principle of relevance
Human cognition tends to be geared to the maximisation of relevance. (Sperber & Wilson, 1986b, p. 260)

'Relevance' here is a technical term. It is defined as a property of inputs to cognitive systems: an input is more relevant the more cognitive effects it yields, and less relevant the more mental effort it takes to process.

Relevance of an input to an individual
a) Other things being equal, the greater the positive cognitive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.
b) Other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time. (Wilson & Sperber, 2004, p. 609; c.f. the original formulation, at Sperber & Wilson, 1986b, p. 153)

On this definition of relevance, the cognitive principle is the claim that human cognitive systems tend to work with their input in such a way as to yield the maximum cognitive benefit for the least mental effort. The reach of this principle is rather broad. For its purposes, cognitive systems include (at least) those that are centrally involved in perception, memory and reasoning as well as those that underpin the production and interpretation of utterances.

The definition of relevance obviously raises two questions: i) what constitutes cognitive effects; and ii) what causes mental effort? Relevance theory gives definite, although not necessarily exhaustive answers to these two questions, and I set them out below.

A less obvious question concerns the cognitive principle: *How* do cognitive systems maximise relevance? Is it, for example, by systematically minimising effort or by systematically maximising benefit? It is compatible with the cognitive principle that different cognitive systems implement different approaches to maximisation. However, we will see below a) that relevance theory has a general account of how the mind as a whole directs effort to tasks that yield cognitive effects, and b) that much more specific claims are made about how the system for interpreting utterances seeks relevance(for discussion see Sperber, 2005). But before I go into these answers, I want to sketch out the intuitive reasons for the core assumptions set out above.

Relevance theory starts from the idea that there is normally much more going on in the environment of any human being than it could pay attention to, and certainly much more than it could mentally process fully. (For discussion, see Sperber & Wilson, 1996; Allott, 2008, ch. 3.) If this were not the case, there would be no need to consider a trade-off between the effort put in and the benefit extracted from doing so. We could process each input fully to extract all the cognitive benefit it might yield, and theories of cognition could ignore processing effort. However it is highly plausible that the environment is too full, and processing too costly, for this abstraction to be justified, particularly considering that by 'environment', here, one must understand not just physical objects, but also sources of information such as utterances made by other human beings, books, the internet, advertisements etc.[6] (Sperber & Wilson, 1996, p. 530; Todd & Gigerenzer, 2000, pp. 729–730). This crucial assumption which underlies relevance theory – that we cannot maximise by considering all options and processing each of them as deeply as possible – is shared with work on 'bounded rationality', pioneered by Herbert Simon, and including

---

6 The 'environment' of each cognitive system is still richer, since it includes outputs from other cognitive systems. For example, our general reasoning is fed by memory, not just by our perceptions of the external environment.

research on 'simple heuristics' by Gerd Gigerenzer and colleagues (Simon, 1957; Cherniak, 1981; Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999). I return to these parallels in the discussion of the relevance theoretic comprehension procedure in section 3.7 below.

A further assumption is required to justify the conclusion that our cognitive systems tend to get a good return on effort expended. That assumption is, roughly, that our cognitive systems are well-adapted to their normal environments. In lectures, Sperber quotes the biologist Dobzhansky: "nothing makes sense in biology except in the light of evolution" (Dobzhansky, 1964, p. 449). Human beings are evolved creatures and complex subsystems including physical organs like the heart, brain and skin and cognitive systems such as memory, face-recognition, ability to communicate etc. must therefore be seen as having been subject to selection pressure. In addition, children's abilities and knowledge develop from infancy, assuming that the child is in an appropriate environment. Thus we should expect 'normal' adults, on average, to be well adapted to normal environments.

There is an analogy with an animal that forages for food, such as a monkey living in the canopy of a rainforest. It will look for things that have a high nutritional payoff: ripe fruit probably contain more energy than leaves, for example. But the monkey cannot just be built to pursue high-energy food at any cost. There must be some balancing of the nutritional payoff against the costs required to obtain and process the food. Fruit that are far away and hard to reach are not as good as fruit that are to hand. A well-adapted creature should tend to eat nearby fruit first, before investigating food that is up at the end of narrow branches and difficult to reach. Equally, we would expect it to go for food that can be eaten straightaway if it can find them, rather than fruit or nuts with hard shells that require a great deal of effort to open. That is not to say that monkeys never bother with fruit that are difficult to process: in fact, some of them use stones to smash open tough fruit, seeds and tubers (Moura & Lee, 2004, p. 1909), but presumably they only do this if the tough food is much more nutritious than the other available food sources[7].

According to relevance theory, something very similar applies to human cognition. The cognitive system should (if it is well adapted) be so constructed that it seeks and processes inputs that are cognitively valuable, all other things being equal; and, on the other hand, that it looks for things that are easy to process, all else being equal. If something is difficult to process, then it will only be worth attending to if the payoff is big enough (where how big that is depends on the other possible sources of cognitive nutrition in the environment, and on the organism's general state of alertness and stores of energy). Conversely, if an input has a low payoff then it will only be worth processing if that is easy to do (where, again, how easy that needs to be depends on the other potential sources of relevance, and alertness and energy).

As well as these parallels with foraging theory, relevance theory's fundamental dependence on notions of cost-benefit trade-off and maximisation make it an intellectual cousin of game theory and rational decision theory, areas which study decision making on the assumption that agents are rational maximisers[8]. The parallel is closer with fields such as foraging theory and evolutionary game theory than with standard game theory (Allott, 2006, p. 147). The basis of the models in these fields (as of the cognitive principle of relevance) is not that agents or their cognitive systems are aware of all the potentially relevant details of the structure of the environment, nor that they use this information to maximise rationally – the 'Common Knowledge and Rationality' assumptions of standard game theory. Insofar as the cognitive principle of relevance is a principle of rational maximisation, the kind of

---

7 Moura and Lee say that the capuchins they studied, "living in a harsh dry habitat, survive food limitation and foraging time constraints through their extensive tool use." (p. 1909). On animal foraging more generally, see Emlen, 1966; Stephens & Krebs, 1986; Stephens, Brown & Ydenberg, 2007.

8 Optimal foraging theory is also in this intellectual territory, since it can be seen as an application of rational decision theory.

rationality involved is of the evolutionary, adaptive sort: that is, it is assumed that evolution and development have selected for systems which produce behaviour that tends to maximise return in normal environments by working with limited information and taking shortcuts.

It should also be clear that no higher-level rationality is necessarily involved, that is, the kind of rationality that requires awareness of and openness to reasons, the ability to reflect on actions and their consequences and so on (Evans & Over, 1996; Sloman, 1996). Of course, human beings are (sometimes) capable of such reflection, but it is not our reflective abilities that are supposed to underwrite the adaptive rationality summarized in the cognitive principle. Rather, the cognitive principle is supposed to apply to all aspects of human cognition, including such largely automatic, non-reflective systems as the face-recognition module and our innate tendency to attend to loud noises, as well as to reflective, conscious, 'person-level' reasoning.

2.2 The payoff: Cognitive effects

> A positive cognitive effect is a worthwhile difference to the individual's representation of the world (Wilson & Sperber, 2004, p. 608).

> … the addition of new information which merely duplicates old information does not count as an improvement; nor does the addition of new information which is entirely unrelated to old information. The sort of effect we are interested in is a interaction between old and new information. (Sperber & Wilson, 1986b, p. 106)

In relevance theory, benefit to cognition is seen as a matter of the positive cognitive effects – the worthwhile changes in the individual's cognitive system, including improvements in her representation of the world – that are produced in an individual by processing an input in a context[9]. Changes in the representation that make it less good for "the fulfilment of cognitive functions or goals" (Sperber & Wilson, 1995, p. 265) (such as changes that take it further away from accurately representing the world) are cognitive effects, but not positive ones, and they contribute not to actual relevance but (in some cases) to how relevant an input *seems* (Sperber & Wilson, 1995, p. 263ff)[10].

What counts as a improvement in an individual's representation of the world? Sperber (2005, p. 65) lists several ways that our knowledge can be fruitfully revised as a result of processing new inputs:

> adding new pieces of knowledge, updating or revising old ones, updating degrees of subjective probability in a way sensitive to new evidence, or merely reorganizing existing knowledge so as to facilitate future use.

Simplifying a bit, the three types of cognitive effect normally discussed in relevance theory are as follows:

Cognitive effects

(1) to support and strengthen an existing assumption;

(2) to contradict and rule out an existing assumption;

---

9 Cognitive effects are sometimes called contextual effects, particularly in Sperber & Wilson, 1986b.

10 This is a change from the definition of cognitive effects in Sperber & Wilson, 1986b. Sperber (2005, p. 65) suggests that in practice "the brain would be roughly right in treating any and every cognitive effect as a positive effect, in other words, as a cognitive benefit."

(3) to interact inferentially with existing assumptions to produce a new conclusion.

In this definition, assumptions are mental representations of aspects of the world: propositions that are believed by the individual, or at least given some degree of credence.

The first kind of cognitive effect is to raise the degree of credence that an individual accords to a particular assumption: e.g. from *probable* to *almost certain.* For example, Mary, who is about to enter King's Cross station, believes that it is probable that there will be a train to Newcastle within the hour (since she believes that there are several each hour during the daytime, and that it is daytime, and has no good reason to think that there is a rail strike, etc.). Entering the station she sees that it is 9.20 and there is a 9.46 train for Newcastle listed on the departure board. Her original belief is strongly reinforced.

The second type of cognitive effect is to reduce to nil the credence that the individual attaches to an assumption. Suppose that when Mary looks at the departure board the first Newcastle train listed is at 11.20, or that the board is displaying a notice saying 'All trains cancelled'. Either of these bits of input would contradict her original belief and – in normal circumstances and absent contrary evidence – either would be credible enough to rule it out[11].

As an illustration of the third type of cognitive effect, suppose now that Mary knows that there is a newspaper shop in the station, and has normal beliefs about how long it takes to buy a newspaper, and, once she is in the station, can see how far it is from the shop to the train. When she enters the station at 9.20 and sees that the next train is at 9.46 she may infer that she has time to buy a newspaper before boarding the train. This is a cognitive effect of the third type. The new input – the time of the next train – interacts inferentially with assumptions that were already available to Mary – about the availability of newspapers, and the time taken to get one and to get to the train – to yield a new conclusion.

Note that it is part of the criterion for this to be a cognitive effect that the interaction between the beliefs is inferential. From *It is 9.20*; *The next train is at 9.46*; and *It takes no more than 10 minutes to buy a newspaper here*, it follows that there is time to buy a newspaper, so this is a *bona fide* cognitive effect. In contrast, an input that causes a new assumption in a purely associative way does not count as a cognitive effect: e.g. the thought that the train is at 9.46 reminds you of granny since she lives at number 46, which in turn reminds you that you should visit her soon. I return below to the assumptions that relevance theory makes about the role of inference in cognition.

Note also that it is a deliberate feature of Sperber and Wilson's characterisation of cognitive effects that learning new information that has no relation to any previously held assumption does not count as a cognitive effect, even if the new information is true. They say that new information that "is entirely unconnected with anything in the individual's representation of the world … can only be added to this representation as isolated bits and pieces, and this usually means too much processing cost for too little benefit." (Sperber & Wilson, 1986b, p. 48)

2.3 The cost: processing effort

What is meant in relevance theory by 'processing effort' is the effort required to process an input *to the point that its cognitive effects are derived*. More specifically, this is the effort taken to "to represent the input, access contextual information and derive any cognitive effects" (Wilson, 2009, p. 394). This effort

---

11 One might wonder why the lowering of credence in an assumption only counts as a cognitive effect if it lowers it to zero i.e. eliminates it as an assumption. Briefly, it is because Sperber and Wilson assume that "[mere] weakening is always a by-product of a more basic contextual effect" (Sperber & Wilson, 1986b, pp. 294, fn d), for example the elimination of another assumption which provided support for the one that is weakened – and so mere weakening does not need to be counted separately.

is therefore a sum of the effort involved in perception, memory and inference (Wilson, 2009, p. 394).

Beyond this general characterisation, relevance theory does not try to define sources of processing effort *a priori*. Instead it works with the results of the fields of psychology which study perception, memory and inference. Relevant research includes work on attention in perception (e.g. Lavie, 1995; Pashler, 1998; Lavie, 2001), in psycholinguistics on retrieval of word senses and disambiguation, which has tended to focus on effort factors, (e.g. Meyer & Schvaneveldt, 1971; Neely, 1991; Forster & Chambers, 1973), and in the psychology of reasoning on the varying costs of different types of inference (e.g. Braine, 1978; Braine & O'Brien, 1998; Rips, 1983; Johnson-Laird, 1983).

Different stimuli will in general require different amounts of processing effort. For example, a longer sentence will (other things being equal) require more effort to process than a shorter one. An uncommon word, or an uncommon sense of an ambiguous word, requires more effort to process than a common one (Forster & Chambers, 1973).

A more subtle point is that the same stimulus in different contexts will generally require different amounts of processing effort. This is because in different contexts the stimulus may be more or less salient (i.e. more or less easy to perceive); the contextual assumptions required to process it may be more or less accessible (i.e. more or less easy to retrieve from memory or derive); the inferences required to draw out its implications may be more or less involved and demanding, and, indeed, what implications it supports will also depend on the context (Wilson & Sperber, 2004, p. 609).

2.4 How do we maximise relevance?

> Within relevance theory, the problem is not so much to assess contextual effects and processing effort from the outside, but to describe how the mind assesses its own achievements and efforts from the inside, and decides as a result to pursue its efforts or reallocate them in different directions. (Sperber & Wilson, 1986b, p. 130)

The picture of cognition that relevance theory assumes is of a number of possible inputs dealt with by a number of mental processes running in parallel. Processes and inputs that are cognitively productive – e.g. returning a lot of effects for reasonable effort, or returning reasonable amounts of effects for low effort – will be preferentially given resources:

> cognitive resources tend to be allocated to the processing of the most relevant inputs available … … … human cognition tends to be geared to the maximisation of the cumulative relevance of the inputs in processes. It does this not by pursuing a long-term policy based on computation of the cumulative relevance achieved over time, but by local arbitrations, aimed at incremental gains, between simultaneously available inputs competing for immediately available resources (Sperber & Wilson, 1995, p. 261)[12].

In a system like this, there is no need for the cognitive systems to calculate ahead of time what the relevance of an input is going to be. That is just as well, since that would probably be self-defeating, requiring huge processing effort (Sperber, 2005, p. 64). The reason is that it is very costly to calculate an optimal stopping point for a search. Simple heuristics that process until some target is achieved, or threshold reached, are much less computationally expensive (Sperber & Wilson, 1986b, pp. 130–131; Todd & Gigerenzer, 2000, pp. 729–730; Gigerenzer, 2004, p. 391; Allott, 2008, pp. 170–172).

In fact, relevance theory takes an even stronger line here. It claims that generally we do not mentally

---

12 See also Sperber, 1994a, pp. 46–50 and Sperber, 2005, p. 63ff.. Similar models include the pandemonium model (Selfridge & Neisser, 1960), and 'enzymatic computation' (Barrett, 2005).

represent processing effort or cognitive effects, so they could not enter into calculations of whether to proceed in processing, and that when represented at all they are represented as comparative (not absolute or quantitative) judgments. Our awareness of mental effort and effects, Sperber and Wilson speculate, may depend on our awareness of "symptomatic physico-chemical changes" that they cause (1986b, p. 130) in much the way that we have a sense of how much physical effort is being taken up in lifting a certain object, or how filling a meal is (see also Sperber, 2005, pp. 64–66).

In support of the assumption that we do not, in general, mentally represent mental effort or effects, Sperber and Wilson argue that we are not in fact able to "compare the contextual effects and processing effort involved in any [arbitrary] pair of mental performances" (1986b, p. 131) and that it is "implausible that human beings might have a system for computing and representing the strength of assumptions which is both wholly unconscious and radically more sophisticated than anything that is reflected in their conscious intuitions" (1986b, p. 79).

On the assumption that effects and effort are not mentally represented, it follows that relevance, which is defined in terms of them, is also a non-representational notion. For Sperber and Wilson, "relevance is a property which need not be represented, let alone computed, in order to be achieved" (1986b, p. 132). As with effort and effects, our sense of relevance is intuitive and comparative, rather than absolute.

To summarize: there are two reasons why the search for relevance cannot be driven by calculations of how profitable it will be to process an input: i) the processing required would be too costly; ii) the quantities required for the calculation are not mentally represented, and therefore are not available to be computed over.

Instead, then, our search for relevance is fed by our internal sense of how cognitively profitable and demanding a certain input or task has been proving. Of course, the allocation of resources should also be guided by expectations of future cost and benefit. Monkeys may anticipate good returns from foraging in a guava tree, and we surely expect more cognitive nutrition from a book by Chomsky than from one by Dan Brown. As Sperber and Wilson put it, as well as *retrospective* intuitions, we have *prospective* intuitions about the effort a task will take and the effects that will be achieved (1986b, p. 130). As we will see, it is central to relevance theory's account of utterance interpretation that in communicative interactions there is a very specific expectation about the degree of relevance that each utterance should attain.

2.5 Cognition, mental representation and inference

The assumptions that relevance theory makes about the way that human cognition performs inference are fundamental to the explanations that it gives of utterance interpretation, and more generally to understanding its definitions of cognitive effects and therefore relevance. They are set out in some detail in chapter 2 of Sperber & Wilson, 1986b, but they have not been much discussed in subsequent work. For the purposes of the brief summary possible here, is is convenient to divide them into two postulates, one of which – the Computational/Representational Theory of Mind (Fodor, 1975) – is adopted in some form across much work in cognitive science, and one of which – the deductive device – is more specific to relevance theory.

The first assumption, then, is that cognition can be modelled in terms of computations performed on mental representations. This assumption is what I am calling the Computational/Representational Theory of Mind (C/RTM). This theory has two central commitments. The first is that the form of a mental representation determines the way that it is processed, since the computational rules that operate on mental representations are sensitive only to their formal (i.e. syntactic) properties. For example, the following representation has the form 'P and Q':

(1) John studies linguistics and Mary studies philosophy.

Starting from any representation with the form 'P and Q' as a premise, one can deduce P as a conclusion: in this case *John studies linguistics*. (Of course, one can also deduce Q as a conclusion.)

The second central commitment of C/RTM is that the mind's syntactic operations generally preserve semantic value. For systems that take propositional input and produce propositional output the value preserved will be *truth*. We can see that this is the case for the example given, because any situation in which the premise (John studies linguistics and Mary studies philosophy) is true is one in which the conclusion (John studies linguistics) is also true.

The point of C/RTM is that it provides some insight into the way that one thought leads to another in reasoning and inference: the syntax of a thought "determine[s] the causes and effects of its tokenings in much the way that the geometry of a key determines what locks it will open" (Fodor, 1985, p. 93). Representations of the form 'P and Q' both lead to (i.e. cause) and logically entail representations of the form 'Q'. Thus, Fodor writes, "the syntactic theory of mental operations provides a reductive account of the *intelligence* of thought." (1985, p. 98. His emphasis.) I do not go into more detail here about C/RTM. It has been discussed at length in the philosophy of mind and, as noted, something along these lines is commonly assumed in cognitive science[13].

Sperber and Wilson's second postulate about the cognitive realisation of inference is that human beings possess a *deductive device.* On their view, human beings are not only equipped with the ability to make logical deductions based on the form of mental representations. In addition, i) these rules are "spontaneously brought to bear in the deductive processing of information" (Sperber & Wilson, 1986b) and ii) this also plays a central role in spontaneous non-demonstrative reasoning. I explain the first of these points here, and return to the second point in section 3.8 below, after discussing the role of non-demonstrative reasoning in utterance interpretation.

The deductive device starts with some input (the premises) and performs all deductive inferences that are possible from each premise and from the premises taken in conjunction, recursively (i.e. also operating on the output of the rules), where what is possible is determined by the set of rules possessed by the deductive device, and by the context in which the input is processed. According to Sperber and Wilson, the deductive device has elimination rules like the one mentioned above that takes input of the form 'P and Q' and returns 'P' as output. So if you start with 'John studies linguistics and Mary studies philosophy' you immediately have available to you 'John studies linguistics'. Suppose now that you already believe, falsely, that 'If anyone studies linguistics then he/she is a polyglot'. In that case you will put that together with what you have just deduced and infer, perhaps wrongly, that John is a polyglot. In the terminology of relevance theory, in this example 'John studies linguistics' is an *analytic consequence* of the input (i.e. one that can be reached purely through the use of elimination rules); 'If someone studies linguistics then he/she is a polyglot' is a *contextual assumption*; and 'John is a polyglot' is a *contextual implication* of the input (i.e. an implication of the input taken together with one or more contextual assumptions). Contextual implications are identical to the third type of cognitive effect in the definition of cognitive effects given in section 2.2 above.

There are a lot of details to spell out about how the deductive device works, and not enough room here to go into them in depth. There are similarities to the system for deduction postulated in 'mental logic' theories of reasoning (e.g. Braine, 1978; Rips, 1983; 1998). One important detail is original to relevance theory. In order to stop overgeneration, Sperber and Wilson postulate that the deductive device does not have introduction rules (1986b, p. 96). For example, given a representation of the form 'P', it

---

13 e.g. Newell & Simon, 1976, who call their version of the framework the 'Physical Symbol System Hypothesis'. For discussion of C/RTM see Barrett, 2005, pp. 259–263; Allott, 2008, p. 105ff.

does not generate a representation of the form 'P or Q', although this would be a logically impeccable deductive inference. Another crucial detail is that the output of the deductive rules is monitored for redundancy (generating something that is already present) and for contradiction (generating a mental representation that is the logical negation of one that is already present), and in each case, suitable action is taken (Sperber & Wilson, 1986b, p. 95).

The deductive device is important for relevance theory in several ways. One important point is that the three types of cognitive effect discussed in section 2.2 above are a corollary of the way that the deductive device is defined. That is, given the specifications of the deductive device, one can show that there must be at least those three types of cognitive effect (1986b, pp. 108–109). So the assumptions made about the deductive device feed into the definition of cognitive effects and therefore into the definition of relevance.

There is another reason why the deductive device is important to relevance theory's account of utterance interpretation. Sperber and Wilson postulate that it plays a central role in non-demonstrative inference (1986b, p. 108). In deduction, if the premises are true and the deductive rule is sound and correctly applied then the conclusion must be true, as with the inference discussed above, from *P and Q* to *P*. Non-demonstrative inference is different, in that it is inference that is uncertain. One type of non-demonstrative inference is inference to the best explanation, in which there is some event or state of affairs and we want to know why that event happened or how the state of affairs came to be. On the Gricean view of communication adopted by relevance theory, interpreting an utterance is just this sort of problem. In the next section I set out the central assumptions that relevance theory makes about communication, and I return to explaining the role of the deductive device in inference to the best explanation once I have shown why utterance interpretation is seen this way.

3 Communication and relevance

Relevance theory may be seen as an attempt to work out in detail one of Grice's central claims: that an essential feature of most human communication, both verbal and non-verbal, is the expression and recognition of intentions. (Wilson & Sperber, 2004, p. 607, referring to Grice, 1989: Essays 1-7, 14, 18; and Retrospective Epilogue).

3.1 The problem: inference about intentions

Turning to the second part of the core of relevance theory, we come to the problem that relevance theory was devised to solve. How do human beings communicate? More specifically: How is it possible that in saying a phrase and/or making gestures, a human being can convey certain propositions to a conspecific? Conversely, how can the conspecific who has perceived the utterance work out what are the propositions that the producer of the utterance had in mind?

The way I have stated these questions already implicitly narrows down the field of phenomena to be explained. We are concerned here with deliberate communication – utterances made on purpose – rather than the sort of accidental information transfer that results from non-deliberate signs or signals: one's accent, posture, pheromones etc. Relevance theory adopts this more precise and narrowly focussed version of the problem from the work of the philosopher Paul Grice (and refines it somewhat for the study of communication, a shift from Grice's interest in 'speaker meaning', as discussed below). In relevance theory, this kind of deliberate, open communication is called 'ostensive-inferential' communication.

Crucially, relevance theory also accepts Grice's characterisation of an utterance as the utterer's

expression of certain intentions. There are two strands of Grice's work that are relevant here: his theory of conversation and his theory of meaning. In the latter work, Grice tried to give a definition of meaning and the verb 'mean' for cases of communicative meaning, or as he called it, *speaker meaning* (excluding another use of the word 'mean' which is typified by such examples as 'Smoke means fire' and 'Black clouds mean rain').

According to Grice, when a speaker means something by an utterance the speaker has a set of nested intentions. The first of these is the intention to produce a certain response in the hearer. In the terms of a cognitive theory, we can think of this as an intention to modify the hearer's mental representation of the world by providing the hearer with information about the speaker's representation of the world. To take a simple case, when a speaker says 'It is sunny', she may intend her addressee to come to think it is sunny. The reason that the hearer comes to think this (if he does) is that the utterance provides *prima facie* evidence that the speaker thinks that it is.

The second intention is that the first intention be recognised. This criterion rules out cases in which an agent wants to bring about a change in the hearer in some other way than by openly producing an utterance. Famously, Grice discusses a case in which Mr X is anonymously informed of his wife's affair by means of a photograph that has been left lying where he will see it. In this case, the person who places the photograph intends to affect Mr X's beliefs (so has the first intention), but does not want Mr X to know that she intended any such change in his beliefs (nor indeed that she had anything at all to do with the photograph), so she lacks the second intention.

This basic structure is adopted by relevance theory as characteristic of ostensive-inferential communication, in the following form:

Ostensive-inferential communication

a) *The informative intention*:
   The intention to inform an audience of something.

b) *The communicative intention*:
   The intention to inform the audience of one's informative intention. (Wilson & Sperber, 2004, p. 611. See also Sperber & Wilson, 1986b, pp. 46–64.)

As demonstrated by Grice's photograph example, the presence of the communicative intention is a criterion for whether the speaker intends to communicate in the deliberate, purposive sense that we are discussing. Moreover, the success of this intention is sufficient for successful communication[14]. That is because if this intention succeeds, then by definition the hearer has recognised the informative intention: i.e. he realises what it is that the speaker intended him to come to think. The success of the informative intention, by contrast, is not required for successful communication. Believing what a speaker has communicated is a different matter from understanding. For example, an utterance of 'It's sunny' is understood when the hearer grasps that the speaker intended him to think that the weather is sunny. Whether he trusts the speaker enough to believe her is a separate matter. (Sperber, Clément et al., 2010 discuss this last point thoroughly).

For Grice, the decision to include this second intention in his theory of speaker meaning rests on intuitions about whether it helps to capture the intended sense of 'mean': e.g. we might be reluctant to

---

14 But not necessary, according to Sperber and Wilson. They suggest that unintentional ostensive communication is possible in cases where an utterer has the informative intention but not the communicative intention (as in the photograph example) but acts so ineptly that the intended audience infers that the informative intention is present (1986b, pp. 63–64).

say that the anonymous photograph-placer *meant* that Mr X was having an affair (or anything at all) by the photograph (or her leaving of it where Mr X would see it).[15] For Sperber and Wilson, whose concern is not conceptual analysis but the foundation of a scientific account of communication, the criterion is different. The hope is that the presence of the informative and communicative intentions marks out a natural class of phenomena ('ostensive stimuli' as they are called in relevance theory) which fall under interesting generalisations and laws and can be productively studied.

Scientific study of any area works towards lawlike generalisations relating to that area. Phenomena that fall under a particular natural law are described as a natural kind[16]. The existence of a general term (such as 'communication') in itself provides no guarantee that there is any such natural kind. Sperber and Wilson give *locomotion* as a counter-example (Sperber & Wilson, 1986b, pp. 2, 3). There is, they say, no general theory of locomotion. There are specific theories of certain modes of locomotion – of aerodynamics, of the biophysics of walking, of flight and of swimming – and more general theories that are relevant, including laws of motion and of the conservation of energy. But there are no interesting, lawlike generalisations that hold at the level of locomotion rather than at a more specific or more general level. In that sense, locomotion is not a natural kind.

What is the evidence that ostensive stimuli form a natural kind? As in all scientific research, the ultimate criterion is the success of the theory which is founded on the assumption, relative to competing research programmes. If sufficient progress is not made, then the conclusion should eventually be drawn that foundational assumptions are wrong (or at best, unproductive).[17]

There are some pre-theoretic intimations that the choice of ostensive stimuli as an area of study will be productive, and considering them leads directly to the next core assumption of relevance theory, the inferential model of communication. Communication seems to require separate study from linguistic syntax and semantics for two reasons, both implicit in Grice's work. As Levinson writes,

> Grice's theory gives us an account both of how we can communicate without conventional signals at all... and of how we can communicate something distinct from what the conventional signals actually mean. (Levinson, 2006, p. 50)

The first point here is that it is intuitively clear that both gesturing and uttering linguistic material are (or rather, can be) means of communicating. Indeed most spoken utterances involve both simultaneously. Crucially, speaker intentions are normally taken as criterial for communicative gestures as well as for linguistic utterances. When we see someone pointing, and want to know what she meant by it, then what we are interested in is finding out what she *intended* to point to. There may be many objects and parts of objects in the direction she pointed in; but what matters is which one she had in mind and wanted her audience to come to have in mind[18]. This is parallel to the case of assigning

---

15 Grice's definition of speaker meaning includes a third intention, expressly to rule out from counting as meaning such cases as openly showing a photograph as evidence of an affair, showing a plaster cast as evidence of a broken leg, or showing a severed head as evidence that the person whose head it was, is dead. While it might be that these are not happily called cases of *meaning* (which as discussed, was Grice's concern), there is no doubt that they are cases of *communication*, in the deliberate, open sense that we have been discussing, so this third intention is not needed in relevance theory (Sperber & Wilson, 1986b, pp. 53–54).

16 The term 'natural kind' comes from Quine, 1969. The criteria for natural-kind-hood are debated. Bird & Tobin, 2010 discuss various criteria.

17 Sperber and Wilson give semiotics as an example of a field that has failed to progress partly because its fundamental assumptions do not pick out a natural class: there are no interesting generalisations, they say, to be obtained over the totality of languages, fashion, novels, road signs etc. seen as coded signals (Sperber & Wilson, 1986b, p. 6ff).

18 What the producer of an utterance intends to communicate is constrained (like other intentions) by what she can rationally hope to achieve. For example, I could not normally expect an addressee to work out that I intend to talk about cats using the word 'dog', nor that I intend to refer to my cat by pointing at a passing dog, so I cannot normally intend these

reference to indexical linguistic items such as pronouns. If a speaker says (for example) "It'll be here later", then questions about what 'it' means in that utterance are really questions about what the speaker intended to refer to. So it seems that we need a theory that covers both non-verbal and verbal communication and relates them both to speaker intentions. Grice's work on speaker meaning provided the basic framework for such a theory.

A further point is that gestures need not have any encoded meaning (in Levinson's terms, there need not be any 'conventional signal'). Sperber and Wilson give the example of raising one's empty glass in a pub, so as to draw a friend's attention to it (2008, p. 89). There is no code or convention that says that raising one's glass means 'Please get me another drink', but in the right circumstances the gesture would be understood as conveying that. Again, concern with the meaning of the gesture on a particular occasion comes down to interest in the utterer's intentions. One might ask the utterer: *What did you mean by raising your empty glass like that?* or *What were you trying to convey?*

Note also that if the speaker had no intention to convey information but was (e.g.) holding up the glass to better examine it in the light, then we would say that the gesture was not a communicative act at all. So the intentions of the maker of the utterance seem to be more fundamental to communication than are language or codes more broadly.

More precisely, as Sperber and Wilson put it:

> Grice's greatest originality was not to suggest that human communication involves the recognition of intentions. That much … is common sense. It was to suggest that this characterisation is sufficient: as long as there is some way of recognising the communicator's intentions, then communication is possible. (Sperber & Wilson, 1986b, p. 25)

Even in the cases of linguistic utterances (and utterances of gestures that encode meaning[19]), the meaning of the utterance may differ from the encoded meaning of the phrase or gesture uttered. Here it is Grice's theory of conversation that is directly relevant. Grice discussed examples in which intuitively what the speaker means includes something that the speaker intentionally implies by (or in) making her utterance: that is, an implicature. Grice intended this category to unify such apparently diverse phenomena as indirect answers to questions and ironic utterances.

(2) Mary: Have you done the hoovering?
    John: I've only just got in from work.

(3) What lovely weather! [said in a downpour]

The intuition that Grice trades on here is that in both cases at least part of what the speaker means is something quite different from what she says. He coined the word 'implicature' as a term of art for this sort of thing: an intended implication of an utterance.

As Grice pointed out, utterances of a given sentence may have different implicatures (or none) on different occasions, in different contexts. Also a speaker may cancel an implicature (e.g. if John's reply in (2) were 'I've only just got in from work, but in that short time, yes, I've already done it', the 'but'-clause cancels the implicature of the previous clause). Implicatures, then, are not something that should be treated by linguists as encoded in the words uttered, but are instead things that the speaker communicates by relying on the hearer's ability to work out that the speaker intended to convey

interpretations (Grice, 1971; Sperber & Wilson, 1986b, p. 169; Neale, 1992, p. 551).

19 Some but not all gestures encode meanings. For example, thumbs-up encodes something like 'Good!' (or, for divers, 'Let's surface').

something distinct from what she said.

To summarize, communication is distinct from linguistic encoding in that it can be accomplished by gestures with no conventional meaning, and in that speakers often communicate something different from what is encoded by the words or gestures they utter. In establishing these points, Grice implied that communication cannot be purely a matter of encoding and decoding (or in more Gricean terms, the deploying and retrieving of 'timeless' meanings of words) and that recognition of speaker intentions is sufficient for communication. But then how does this work?

In his theory of conversation, Grice outlines a way in which implicatures could be inferred by hearers, and therefore that speakers can rationally intend to convey them. The details of Grice's theory do not matter here (but see remarks below on contrasts with relevance theory's communicative principle of relevance). What is crucial is that his theory of conversation proposes that hearers must *infer* what it is that speakers intend to convey. Thus, as Wilson and Sperber put it, "Grice laid the foundations for an inferential model of communication, an alternative to the classical code model." (2004, p. 607)

3.2 The inferential model and the code model

Work on communication in relevance theory is a thorough exploration of the view that the linguistic material in an utterance serves as a clue that the speaker offers the hearer about her communicative and informative intentions and that the hearer uses this evidence to infer an appropriate interpretation of the speaker's utterance. This is a radical departure from pre-Gricean accounts of communication, which effectively treat communication as purely a matter of coding and decoding of a message.

It is worth noting two ways in which the inferential model differs from a code model. First, the code model is a one-stage model of utterance interpretation. The hearer just decodes the signal and retrieves the message. In contrast, in Sperber and Wilson's inferential model there will often be two stages to utterance interpretation, since the clues provided by the speaker will often include a phrase of some language, and language is a code. In such cases the hearer will have to i) decode the phrase used, and ii) infer what the speaker intended to convey by using it[20].

Secondly, according to the inferential model, communication is fallible (and therefore risky but also creative) in ways that purely coded communication is not[21]. The code model claims that where the code is shared by speaker and hearer and the encoded message is received intact, the message will be decoded precisely as it was sent (Sperber, 1994b). On this model, the norm is a kind of reproduction of the speaker's thought in the mind of the hearer.

Contrast this with the inferential model. The kind of inference involved in utterance interpretation is inference to the best explanation. The hearer has to infer intentions that the speaker had and which led her to make the utterance. The input to the inference is something like (e.g.): Mary said: "John isn't here yet" (with a certain intonation, perhaps accompanied by certain gestures). The question, then, is *What best explains the production of these words and gestures (at this time, in this way)?* The answer will generally be of the form, *Mary wanted to convey Σ*, where Σ is the hearer's best estimate of the intended interpretation[22]. As discussed above, this sort of inference is unlike logical deduction in that the explanation reached is not guaranteed to be the right one.

As also noted, the inferential model allows room for creativity. The speaker may leave open to some

---

20 It does not follow that in interpreting an utterance all the decoding is done first, followed by the pragmatic inference. As is well known from psycholinguistics, processing of utterances proceeds 'online', that is, in real time, as the words are heard or read.

21 This is also a Gricean observation. Grice lists indeterminacy as a property of implicatures, a conequence of the fact that they must be inferred non-demonstratively (Grice, 1975, p. 58).

22 I use 'Σ' (for 'sum') because in general an interpretation is a bundle of propositions. See sections 3.4 and 3.5 below.

extent just what she wants to convey in making a certain utterance. Then the hearer will have to take on some responsibility for the interpretation that he derives. Relevance theory's treatment of this point makes use of two related notions introduced by Sperber and Wilson: manifestness and strength of communication.

3.3 Manifestness and strong and weak communication

Roughly, an assumption is manifest to an individual in a context if he could represent the assumption mentally (on the basis of memory, perception or inference) and accept it as true or probably true. Some assumptions are not manifest at all, while among some assumptions that are manifest some are more highly manifest than others. That is, manifestness is both a classificatory and a gradable notion. (Compare, e.g. *poisonousness*: substances may be anywhere from mildly to highly poisonous, or not poisonous at all[23].) According to relevance theory, utterances do not necessarily make the addressee mentally entertain the assumptions communicated. Rather they make it manifest that the speaker intended to make them manifest (Sperber & Wilson, 1986b, p. 197), and the degree to which they are made manifest is variable. Some implicatures are so highly manifest that the speaker will almost certainly recover them – and the utterance would not be relevant enough without them. These are strong implicatures. Others may be only weakly manifest: the utterance provides the hearer with some evidence that the speaker intended to convey them but that evidence is not conclusive, and the relevance of the utterance does not depend on any particular one of them. These are called weak implicatures[24]. Consider (4) (Sperber & Wilson, 1986b, p. 194):

(4) Peter: Would you drive a Mercedes?
   Mary: I wouldn't drive ANY expensive car.

Mary's utterance strongly implicates (5) and (6). If Peter does not grasp them then he has not understood the utterance.

(5) A Mercedes is an expensive car.
(6) Mary wouldn't drive a Mercedes.

In addition it has a number of weaker implicatures, including, from stronger to weaker, (7) a–c. These contribute to the relevance of the utterance, but Peter need not entertain any particular one of these to get the point Mary is making.

(7) a) Mary wouldn't drive a Rolls-Royce.

   b) Mary wouldn't drive a Lexus.

   c) Mary wouldn't drive a Saab.

3.4 Implicated premises and implicated conclusions

---

23 Example suggested by Deirdre Wilson (p.c.). See also Sperber and Wilson (1986b, pp. 39, 79–80).

24 Relevance theory also distinguishes between i) the strength with which an utterance *implicates* an assumption, and ii) the strength with which an assumption is *implied* (Sperber & Wilson, 2008, §7), but there is no space in the current paper to discuss this distinction.

> We assume that a crucial step in the processing of new information, and in particular of verbally communicated information, is to combine it with an adequately selected set of background assumptions – which then constitutes the context… (Sperber & Wilson, 1986b, pp. 137–138)

Mary's utterance in (4) exemplifies another important relevance-theoretic distinction. According to Sperber and Wilson, it is possible to implicate contextual assumptions such as (5), as well as contextual implications of the utterance, such as (6)[25]. The latter are *implicated conclusions*; implicated contextual assumptions are called *implicated premises*. According to relevance theory, all implicatures are of one of these two types (Sperber & Wilson, 1986b, pp. 194–195).

Why assume that some contextual assumptions are communicated? Consider the example again. Given (5), but not otherwise, it follows from what Mary asserts that she won't drive a Mercedes, i.e. (6). The assumption in (5) has to be supplied to make sense of her utterance, so Mary must have intended to make both (5) and (6) manifest.

This illustrates the key point about implicated premises and conclusions. They are tightly related to each other in the following way: given the constructed context, i.e. the implicated premises, the explicit meaning of the utterance logically warrants the implicated conclusions. The complete interpretation of the utterance is thus a logically coherent package. I return to this point in section 3.8 below. Now I turn to relevance theory's treatment of the explicit content of utterances.

3.5 Explicatures: basic- and higher-level

Since Grice's work there has been a gradual understanding that the role of pragmatic inference goes beyond the derivation of implicatures to other aspects of what is conveyed by an utterance, particularly the proposition expressed. How far to take this has been controversial (see Hall, this volume). Relevance theorists have been instrumental in this development (Wilson & Sperber, 1981; Sperber & Wilson, 1986b, ch 4; Carston, 1988; Carston, 2002).

As discussed above, the relevance theoretic position is that any linguistic material uttered is no more than a clue to the interpretation. It follows that not just implicatures, but what is explicitly conveyed by an utterance is pragmatically inferred. Consider an utterance of the sentence in (8). It may be used to express (and in this case, assert) the proposition in (9). Relevance theory calls the proposition expressed an *explicature* (by analogy to 'implicature'). It is an inferential fleshing out of the encoded logical form of the utterance. This fleshing out in general may include reference assignment for indexical elements (e.g. 'I' → Peter; 'it' → the car), disambiguation of ambiguous words or phrases, and enrichment (e.g. 'ready' → ready for the trip to the seaside).

(8) Peter: I'll get it ready in time.

(9) Peter will get the car ready for the trip to the seaside in time to set off early enough to get there by noon.

Peter's utterance may be intended as a promise. In that case it also conveys (10):

(10) Peter promises that he will get the car ready for the trip to the seaside in time to set off early enough to get there by noon.

This is also an explicature of the utterance, given the definition of explicature:

---

25 The notions of contextual assumption and contextual implication were introduced in section 2.5 above.

Explicature

> An assumption communicated by an utterance is an explicature if and only if it is a development of a logical form encoded by that utterance (Sperber & Wilson, 1986b, p. 182. See also Carston, 2002, pp. 116–125.)

The proposition in (9) is the basic-level explicature of the utterance; the one in (10) is a higher-level explicature. Higher-level explicatures are embeddings of the basic-level explicature under speech-act descriptors like 'promise that' and 'ask whether', or attitudinal ones such as 'regret that' or 'be pleased that'.

According to relevance theory, the explicatures of the utterance, like the implicatures, may be strongly or weakly communicated, since here also communication is a matter of making it manifest that the speaker wanted to make an assumption manifest. Putting all of this together, we see that in relevance theory an interpretation of an utterance is in general a bundle of propositions – basic- and higher-level explicatures, plus implicated premises and implicated conclusions – each of which the speaker communicates more or less strongly.

Now I turn to relevance theory's explanation of how the hearer arrives at the interpretation.

3.6 The communicative principle of relevance

According to relevance theory, the search for the correct interpretation of each utterance is guided by "the expectation that utterances should meet certain standards" (Wilson, 2009, p. 393). This idea originates with Grice, although the way relevance theory develops it is quite different. Grice proposed that conversation is governed by a Cooperative Principle and a number of conversational maxims: do not say things that are false; provide enough but not too much information; be relevant; etc. From the hearer's point of view, these can be seen as expectations: that the speaker will be cooperative, where that includes trying to tell the truth, to provide an appropriate amount of information, to be relevant, and so on.

Relevance theory postulates instead that each utterance raises an expectation that it will be optimally relevant. This is because each utterance is an ostensive stimulus, that is, an open attempt to take up some of the hearer's precious attention. This is stated in the communicative principle:

The communicative principle of relevance:
> Every act of ostensive communication communicates[26] a presumption of its own optimal relevance. (Sperber & Wilson, 1995, p. 260)

The presumption of optimal relevance has two clauses, as follows:

The presumption of optimal relevance:

> The utterance is presumed to be
>
> (1) at least relevant enough to be worth the speaker's effort to process it and
>
> (2) the most relevant one that is compatible with the speaker's abilities and preferences. (Sperber &

---

26 Strictly speaking, the claim is that every act of ostensive communication *makes (more) manifest* a presumption of its own optimal relevance.

Wilson, 1995, p. 270)[27]

According to relevance theory, the cognitive and communicative principles are not mentally represented by speakers or hearers nor communicated. They thus have a different status from Grice's Cooperative Principle and maxims, which are supposed to guide behaviour through the awareness of the speaker and hearer that they are in force and should be obeyed[28]. The cognitive and communicative principles are intended to be purely descriptive generalisations, like the laws of physics or biology. The claim is that speakers and hearers conform to these principles without awareness of them and without intending to: "Communicators and audience need no more know the [communicative] principle of relevance to communicate than they need to know the principles of genetics to reproduce." (Sperber & Wilson, 1987, p. 704)

The *general* presumption of optimal relevance has the purely descriptive, non-represented status that the cognitive and communicative principles have. It is the specific presumption that comes with each utterance that, according to relevance theory, is communicated. Relevance theory does not claim that the presumption will always be true, nor that it is always taken as true. According to Sperber and Wilson, "It is enough that the presumption of relevance should be communicated – and it always is – to fulfil its most important role: determining the interpretation of the ostensive stimulus." (Sperber & Wilson, 1987, p. 704)

How does the presumption of optimal relevance help the hearer to infer the correct interpretation? The two clauses set a lower bound and a higher point respectively for the relevance that the hearer is entitled to. The first clause sets the lower bound. It might seem that this lower bound is not well-specified, or is uninterestingly low. What level of relevance is "enough to be worth the speaker's effort to process" the utterance? If we recall the discussion of the assumptions around the cognitive principle, we see that the lower bound is largely set by the environment. The point of the first clause is that an utterance must be worth attending to amid the other possible sources of cognitive effects in the hearer's environment. This may be quite a high degree of relevance, given the limits on human attention, and the fact that other potential sources of relevance may be (or seem) highly relevant.

The second clause, by contrast, strikes many people as too strong. Why should the hearer be entitled to expect the *most* relevant utterance that the speaker is willing and able to provide? Sperber and Wilson explain this in terms of two thought experiments (1995, pp. 268–269; see also Wilson & Sperber, 2002, p. 604). Suppose that a speaker wants her utterance to produce certain cognitive effects in the hearer. Now suppose that there are several possible utterances that she could make that would produce these cognitive effects. One of them would produce just the desired cognitive effects, while the others would produce these plus other cognitive effects. Which utterance should the speaker produce?

Now suppose instead that the speaker has a choice between utterances which would all produce only the desired cognitive effects, but some of which are easier for the hearer to process. Again, the question is: Which utterance should the speaker produce?

The general answer, according to Sperber and Wilson is that "She should choose the utterance that would be (or seem) the most relevant to the addressee" (1995, p. 269). Why? Well, minimizing the hearer's effort is good policy because it makes it more likely that the hearer will pay attention and fully process the utterance, i.e. more likely that the cognitive effects which the speaker wants to cause will occur. And maximising the hearer's returns is also good policy because giving the hearer more information that is of interest to him will also maximise the chances of his paying attention, fully

---

27 The formulation of the presumption of optimal relevance given here is stronger than the one originally presented in Sperber & Wilson, 1986b. See section 5 below.

28 The point of calling the maxims 'maxims' is to suggest that, like Kant's maxims, they motivate agents' actions.

processing the utterance, and remembering the information that the speaker wanted to convey.

Note that none of this assumes anything like Grice's Cooperative Principle. It just follows from the speaker's desire to be understood, and the assumption that the cognitive principle of relevance applies to the hearer: i.e. that the hearer generally seeks maximal relevance. According to relevance theory, speakers exploit that tendency.

Since the communicative principle and presumption of optimal relevance set bounds on the relevance to be expected from any utterance, they obviously help to make tractable the search for an interpretation of an utterance. In fact, relevance theory claims that they licence a specific interpretation procedure, the relevance theoretic comprehension procedure.

3.7 The relevance theoretic comprehension procedure

> if there is one conclusion to be drawn from work in artificial intelligence, it is that most cognitive processes are so complex that they must be modelled in terms of heuristics rather than failsafe algorithms. We assume, then, that communication is governed by a less than perfect heuristic. (Sperber & Wilson, 1986b, p. 45)

The relevance theoretic comprehension procedure is as follows:

(1) Following a least effort path, construct a (hypothetical) interpretation of the utterance. This interpretation will generally include explicatures, implicated premises and implicated conclusions.
(2) Check to see whether the interpretation as a whole satisfies both clauses of the presumption of optimal relevance. That is, it should be i) relevant enough (i.e. it should provide enough cognitive effects for the effort expended thus far in processing the utterance) and ii) it should be the most relevant one that is compatible with the speaker's abilities and preferences (in the hearer's estimation of those abilities and preferences).
(3) If the interpretation hypothesised in step 1 passes the test outlined in step 2, accept it as the intended interpretation.
(4) If not, go back to step 1, and work through the steps again: i.e. construct the next most accessible interpretation and check it for optimal relevance. Repeat until an optimally relevant interpretation is found. Alternatively, or in addition, adjust the expectation of relevance: perhaps the utterance is not as relevant as it might have been because the speaker is not fully competent; or the speaker is not benevolent and the utterance is only intended to seem optimally relevant.[29] If no interpretation that is optimally relevant (or intended to be optimally relevant, or to seem so) is found within reasonable time/effort, the overall cognitive economy will shut down the search.

Note first that the most accessible interpretation will always be checked first, given that this procedure follows a least effort path. So if the most accessible interpretation is relevant enough, it will be accepted as the intended interpretation (i.e. the speaker's intended interpretation). Thus in situations where other sources of potential relevance are largely absent, and where the hearer has low expectations of the speaker's abilities and preferences, the most accessible interpretation will generally be accepted as the correct one[30]. In other cases, the hearer's expectations will make it clear that the speaker intended a more relevant interpretation and the search will continue.

---

29 See Sperber, 1994b, who suggests that the ability to make these adjustments develops in early childhood.
30 As Sperber and Wilson point out (1986b, p. 185), this may be part of the reason why much psycholinguistic work (e.g. on disambiguation) has tended to focus only on accessibility (i.e. effort) factors.

The comprehension procedure is a heuristic in the sense used in the literature on bounded rationality, namely that it is not guaranteed to arrive at the right answer. Like other heuristics that are worth using, it is supposed to find results quickly and without too much effort: it is 'fast and frugal' in Gigerenzer's terms (e.g. Gigerenzer & Goldstein, 1996). In common with other fast and frugal heuristics, it uses both blades of Herbert Simon's scissors: that is, both "the structure of task environments and the computational capabilities of the actor" (Simon, 1990, p. 7). Full computation of all the possible interpretations of an utterance is not necessary, according to relevance theory, because of the environmental regularity described in the presumption of optimal relevance. Instead, because all ostensive stimuli come with a (fallible) guarantee, the hearer's utterance interpretation system just has to come up with the first interpretation that occurs to it, evaluate it, and then perhaps formulate the next most accessible interpretation (which is likely to be a modified variant of the previous interpretation), evaluate that, and so on. In other words, what makes this procedure frugal is that a) it follows a least-effort path, and b) that the first interpretation reached that satisfies the presumption of relevance stops the search.

The reason that the presumption of optimal relevance makes it reasonable for interpretation to follow a least effort path is that relevance varies inversely with effort, so an utterance whose intended interpretation is off the least effort path is less relevant than another utterance that the speaker could have managed to produce. To satisfy clause b of the presumption, speakers have to make their utterances as easy to understand as possible[31]. The reason why the hearer can stop at the first optimally relevant interpretation is that an utterance that has two significantly different interpretations that both yield the expected degree of cognitive effects would fail to be optimally relevant, since the hearer would have to expend effort in choosing between them[32].

3.8 Utterance interpretation as inference to the best explanation

We have seen how the relevance theoretic comprehension procedure decides how to stop and accept an interpretation as the correct, intended one. But how are hypothetical interpretations derived? Part of the answer has already been sketched out above. Verbal material in the utterance is decoded, and then hypothetically fleshed out in ways that include disambiguation of ambiguous words or structure, and the assignment of reference to indexical expressions such as pronouns. This derivation of explicatures occurs in parallel, and in 'mutual adjustment' with the derivation of implicated premises and implicated conclusions. The mechanism for the derivation of implicated conclusions was discussed in section 2.5 above: the deductive device, given an input and contextual assumptions, will churn out contextual implications.

Where, though, do the contextual assumptions come from? A hearer has a lot of information available, some highly accessible, some less so. A linguistic utterance raises the accessibility of information associated with the concepts that are encoded by the words used.

Consider (4) again. Peter expects that Mary's utterance is optimally relevant, and that it will answer his question. Mary's use of the phrase 'expensive car' temporarily makes what Peter knows about expensive cars highly accessible. Putting this together with the fact that he is already thinking about

---

31 For more discussion of why the presumption of optimal relevance mandates a least effort path and stopping at the first optimally relevant interpretation see Sperber & Wilson, 1995, p. 272; Wilson & Sperber, 2002, p. 605; Allott, 2008, pp. 259–260.

32 This is not meant to rule out puns, *double lecture* and the like. In such cases, Wilson and Sperber say, "it is the fact that the speaker has produced such an utterance that is seen as a communicative act. It receives a higher-order interpretation, which may involve endorsing both lower-order interpretations (if they are compatible), or rejecting both (if they are not)." (2002, pp. 605, fn 6)

Mercedes cars, the proposition that they are expensive is highly accessible. At the same time, Mary's utterance is decoded, and reference is assigned to 'I', yielding the proposition: *Mary would not drive any expensive car*. The deductive device automatically combines this hypothetical explicature with the highly accessible *Mercedes are expensive cars* to yield (6), a conclusion that answers his question. Thus there is a logically coherent package of proposition expressed, implicated premise and implicated conclusion. Taking into account also the weak implicatures noted in the discussion of this example above, the total hypothetical interpretation is optimally relevant, so it is accepted as the intended one.

In this model of inference to the best explanation, the inference process is separated into hypothesis formation and hypothesis testing. The hypothesis formation is entirely mechanical. It is partly a matter of taking the most accessible assumptions, the most accessible disambiguation, the most accessible reference assignment etc. and partly a matter of feeding what results through the deductive device. Hypothesis testing is then just a matter of checking the putative interpretation against expectations of relevance[33].

4 Beyond the core

Given that the focus of most research in relevance theory has been on communication, one might wonder whether the broader commitments about cognition (which themselves rest, as discussed, on assumptions about evolution) are strictly necessary.

One way to think about this is to imagine constructing a different theory, which we can call RT′. RT′ adopts the communicative principle and the other assumptions in section 3, but discards the cognitive principle, and with it the evolutionary backstory. The other assumptions in section 2 are kept, including the definition of relevance and the computational/representational theory of mind. The scope of RT′ would be much more restricted. But one can ask what would be lost, from the more specialised perspective of pragmatics, in moving from relevance theory to RT′. A partial answer is that RT′ would lack several important explanatory features. First, and most obviously, it would lack a rationale for the communicative principle of relevance. Why should hearers assume that speakers will be, or try to be, optimally relevant, if there is no general tendency of cognition to maximise returns for effort?

In addition, the cognitive principle implies that humans are somewhat predictable in their cognition, and thus helps to explain how speakers can produce utterances which the hearer will process in the way that was intended.

A related point is that RT′ would have no explanation of why information tends to be stored in long-term memory in ways that are useful for understanding utterances (and for cognition more broadly). The cognitive principle implies that memory should tend to be organised so that information is stored in a useful form and so that it will tend be recalled when relevant, and not otherwise. As discussed in section 3.8 above, the accessibility of information plays a considerable role in relevance theory's explanation of utterance interpretation. It is commonly assumed that long-term memory is organized in chunks sometimes called 'frames' or 'schemas' (Sperber & Wilson, 1986b, p. 138). Thus, for example, when a restaurant is mentioned, it raises the accessibility of stereotypical information about restaurants, such as facts about waiters. The necessity of this sort of structured memory for utterance interpretation is brought out by 'bridging' cases such as the utterance in (11) (c.f. Wilson & Matsui, 1998).

(11) We went to a Thai restaurant. The waiter was from Bangkok.

The cognitive principle provides some explanation why memory is arranged in chunks, and also, therefore, sheds light on how speakers are able to fine-tune their utterances to rely on and exploit such

---

33 See also the more detailed worked example at Wilson & Sperber, 2002, p. 607ff. and discussion at Allott, 2008, pp. 65–66.

facts about what the hearer is likely to have stored and to quickly retrieve.

4.1 Auxiliary assumptions and positive heuristic

I turn finally to some of relevance theory's auxiliary hypotheses and to its 'positive heuristic'. As noted in the introduction, Lakatos proposed that research programmes (or series of research programmes) have positive heuristics, strategies for forming theories outside of the hard core which specify what 'paths of research' to pursue (Lakatos, 1968, p. 168).

> … the positive heuristic consists of a partially articulated set of suggestions or hints on how to develop the 'refutable variants' of the research-programme, how to modify, sophisticate, the protective belt. (Lakatos, 1968, p. 171)

The main thrust of relevance theory's positive heuristic is, of course, something like this: *confronted with a phenomenon(/data) in the realm of communicative behaviour, try to understand it in terms of the operation of the relevance theoretic comprehension procedure, i.e. in a way that is compatible with, and, to as great an extent as possible, predicted by the communicative and cognitive principles*. However, the relevance theoretic research programme has several additional resources that amount to suggestions on how to form theories. I discuss four. Three of these are parts of the framework that are somewhat logically independent of the core, but which are general in that each underlies several auxiliary hypotheses[34]. They are i) the possibility of dividing what is linguistically encoded between conceptual and procedural information; ii) the interpretive/descriptive distinction; iii) the postulation of *ad hoc* concepts. The fourth is not itself a hypothesis, but something more like an attitude to pragmatic theorising: a strong economy principle, akin to Grice's Modified Occam's Razor, but with a wider scope. I discuss this first.

4.2 Modified Occam's Razor

Grice's Modified Occam's Razor is the principle that senses should not be multiplied beyond necessity (Grice, 1989, p. 47). It amounts to an economy argument in favour of treating meanings as pragmatically derived – and thus for Grice, as implicatures – rather than linguistically encoded, unless there is compelling evidence of linguistic ambiguity. For example, an utterance of the sentence in (12) will typically convey that John's kicking of the dog came after his being slapped by Mary (and perhaps also that it was a result of his being slapped). But 'and' does not always convey temporal or causal relation, as illustrated by (13).

(12) Mary slapped John and he kicked the dog.
(13) Mary lives in London and John lives in Oxford.

Grice argued that the word 'and' has just the meaning that it contributes to examples like (13), namely logical conjunction, and that what is additionally conveyed by its use in (12) is pragmatically implicated (Grice, 1967).

Relevance theorists also invoke Modified Occam's Razor. They have been concerned with some of the same words and examples. For example, in a series of papers, Robyn Carston defends the Gricean simple, univocal semantics for 'and' against various objections (Carston, 1988; Carston, 1993; Carston, 2002, ch.

---

34 It is tempting to say that in addition to the core, there is both a *mantle* and a *crust*. Then *ad hoc* concepts, the conceptual/procedural distinction and the distinction between interpretive and descriptive use are in the mantle, while the relevance theoretic accounts of pronouns, utterance modifiers, irony, non-declaratives, loose use, hyperbole and metaphor are parts of the crust.

3; Blakemore & Carston, 2005). The most notable of these is L. J. Cohen's observation that the extra component of meaning cannot be an implicature because it comes under the scope of logical operators (Cohen, 1971). Otherwise the following would seem nonsensical or internally contradictory:

(14) You are being unfair to Mary. It's not true that she slapped John and he kicked the dog. He kicked the dog and she slapped him.

To defend the view that 'and' encodes logical conjunction, Carston takes the non-Gricean position that the pragmatic enrichment in such cases affects the proposition expressed.

Similar use of Modified Occam's Razor is widespread in the work of relevance theorists. The principle is that if something *can* plausibly be done by the pragmatic mechanism, then it *should* be attributed to it, since the pragmatic mechanism is required independently of the analysis of any particular case: it comes for free, one might say. There are real cases of ambiguity which cannot plausibly be analysed as purely pragmatic differences: the lexical ambiguity of 'bank', for example. But note that the pragmatic mechanism will still be required and involved in any genuinely linguistically ambiguous cases, since the hearer has to infer which of the senses the speaker intended.

Relevance theory's use of Modified Occam's Razor clashes with the central positive heuristic of linguistic formal semantics, which is roughly: when one finds a difference in truth-conditions, one should try to show how that difference can be derived compositionally from the encoded meanings of the words in the sentence, postulating complex encoded meanings as necessary. Following this principle leads in the opposite direction to Modified Occam's Razor, locating the explanatory action in syntax and/or semantics and tending to multiply linguistic representations. These are modern variants of Posner's "two competing strategies for the description of verbal communication" (Posner, 1980, p. 170). From the point of view of relevance theory, the issue cannot be settled globally, but only case by case, subject to the Gricean presumption that unless there is good reason to propose two or more linguistic representations, one should prefer a pragmatic explanation.

This way of putting it illustrates that the economy principle of Modified Occam's Razor can be stated in terms of representations. For example, an utterance of the sentence in (12) has representations on at least two cognitively significant levels. There will be a linguistic semantic level (sometimes called LF), and the level of the interpretation of the utterance, after[35] pragmatic processing. That the sentence has (at least) two readings is common ground between relevance theorists and ambiguity-theorists, so there is no clear difference between the theories at the post-pragmatic level. But if we postulate that the sentence is linguistically ambiguous, then the string in (12) will correspond to at least two distinct representations at LF[36], whereas on the Gricean or relevance theoretic position there will be only one. Simply counting the representations we can see that the pragmatic explanation is more economical. And as Carston shows, 'and' has many more than two readings, so the pragmatic account is very much more economical than the proliferation of senses and representations which would be required by a syntacto-semantic account (Carston, 1993, pp. 27–8, 35).

The impulse to simplify carries over, in relevance theory, to whole categories of mental representation, and beyond. As discussed above, relevance theory has only one communicative principle where Grice had the Cooperative Principle and several maxims, and the communicative principle is not mentally represented (except by theorists, of course) whereas in the Gricean framework, speakers and hearers must be aware of the maxims. Then there is also only one way of deriving pragmatically inferred meaning,

---

35 It is from the point of view of interpretation that this mental representation of what is conveyed by the utterance is *after* pragmatic processing: i.e. the hearer only gets to it once pragmatic processing has been performed.

36 One might be able to argue that there is only one LF for the sentence if one postulates (as Cohen does) a very complicated lexical entry for 'and', but this shifts the complexity without reducing it and also creates new problems (Carston, 1993, p. 35).

again in deliberate contrast to Grice's theory, which claims that implicatures arise in at least four ways[37].

In addition, relevance theory rejects both conventional implicatures (that is, implicatures encoded by certain words) and generalised conversational implicatures (implicatures that are pragmatically implied by default). For relevance theory, all implicatures are of the type that Grice called particularised conversational implicatures, the kind that hearers have to a) infer, b) taking into account the specifics of the situation. That is, there is no separate class of default implicatures, and there are no implicatures encoded by lexical items. Finally, relevance theorists do not employ a separate category of presuppositions[38].

4.3 The conceptual/procedural distinction

> Linguistic decoding provides input to the inferential phase of comprehension; inferential comprehension involves the construction and manipulation of conceptual representations. An utterance can thus be expected to encode two basic types of information: representational and computational, or conceptual and procedural – that is, information about the representations to be manipulated, and information about how to manipulate them. (Wilson & Sperber, 1993, p. 1)

According to relevance theory, words can encode two different types of meaning. The first is conceptual meaning. For example, the word 'cat' encodes the concept CAT, and contributes this concept to the proposition expressed by utterances of sentences containing the word, as in (15).

(15) Her cat is antisocial. So no one picks him up and pets him.

Words can also encode procedural meaning, which is to say that they can encode constraints on the way that an utterance is processed. For example, the discourse connective 'so' in (15) encourages a reading in which the second sentence is taken as a conclusion supported by the first. (Contrast the way that 'after all' promotes a reading in which the first sentence is supported by the second.)

(16) Her cat is antisocial. After all no one picks him up and pets him.

The idea of procedural meaning and the treatment of discourse connectives in these terms comes from Diane Blakemore (1987). Subsequently Wilson and Sperber (1993) expanded the role of procedural meaning. They propose that procedural meaning can constrain the derivation of explicatures (whereas previously it had been seen as contributing only to non-truth-conditional aspects of interpretation). For example pronouns are taken in relevance theory to encode constraints on explicatures: e.g. 'him' encodes (roughly) *search for a male individual to fill this slot in the proposition expressed* [39]. So the conceptual/procedural distinction allows a surprising partial unification of the semantics of pronouns and discourse connectives. (For a recent review of work on procedural meaning see the papers in Escandell-Vidal, Leonetti & Ahern, 2011, particularly Wilson, 2011).

---

37 Namely i) real and ii) apparent violation of maxims, iii) clashes between maxims and even iv) in cases in which there is no violation and no appearance of it.

38 Much of what would be explained in terms of presuppositions by other theorists is naturally understood in relevance theory as (and unified with) the communication of implicated premises. Some other alleged presuppositions are treated as entailments: e.g. relevance theorists tend to endorse the Russellian/Gricean account of definite descriptions (Carston, 2002, pp. 110, 306–11). The roots of relevance theory's view of presuppositions are in Wilson's early work (Wilson, 1975b; Wilson, 1975a). See also Kempson, 1975.

39 This account of pronouns, as Sperber and Wilson point out, is something like a cognitive version of Kaplan's character/content distinction.

4.4 The interpretive/descriptive distinction

Relevance theory makes a distinction between different ways that sentences (and words) can be used. Consider examples (17) and (18).

(17) John: What did the prime minister say?
   Mary: He knew nothing about the leak until this week. But I don't believe him. (cf Sperber & Wilson, 1986b, p. 228, their example 101.)

(18) Mary: He's just trying to get himself out of trouble.

An utterance of a sentence can be intended to convey an explicature which is a statement about the world, on the basis that the logical form of the proposition expressed resembles the logical form of a proposition that describes a certain state of affairs. Relevance theory calls this *descriptive use*. Example (18) is most likely to be interpreted this way, i.e. as Mary voicing her own opinion that the prime minister is trying to get out of trouble.

   But this is not the only possible use of a sentence. A sentence can be uttered with the intention that its logical form resembles the logical form of a proposition someone is entertaining, or the logical form of an utterance that someone else has made or might make. In this first case the utterance is a representation of a mental state; in the second it is a representation of another utterance. Relevance theory calls such uses of sentences interpretive, because in typical examples like Mary's utterance in (17) the speaker is not presenting her own view of the way things are, but is acting as an interpreter of someone else's view or utterance[40] (Sperber & Wilson, 1986b, p. 224ff.).

   This distinction is made use of in relevance theory's account of irony. Sperber and Wilson reject the classical definition of verbal irony – the expression of a meaning by using words that usually mean the opposite – and propose instead that irony is interpretive use with a dissociative attitude (Sperber & Wilson, 1986b, p. 237ff; Wilson & Sperber, 1992). In true irony, the dissociative attitude and the fact that there is interpretive use are both tacit – i.e. not linguistically signalled, but left for the hearer to work out. To see what this amounts to, consider the sentences in (19), uttered in a rainstorm. What they would express is similar, but the first is ironic because the hearer is left to infer the attitude and the fact that the utterance is meant interpretively.

(19) a) What beautiful weather!
   b) It was ridiculous to expect beautiful weather.

This account fits well with the observation often made that irony comes at a significant risk of misunderstanding, since on this account ironic utterances require the hearer to infer two pieces of tacitly conveyed information. It also accounts for irony's affinity with quotation. Furthermore, it correctly predicts intuitions for examples that are problematic for the classical definition of verbal irony. Consider (20) uttered in a context in which it is obvious that the car in question has a broken window:

(20) Look, that car has all its windows intact. (Grice, 1967, p. 53)

This is not generally ironic, as Grice noted. The relevance theoretic account predicts this, since it is hard to process as a (mocking) echo of something someone might say or think. However, as the relevance

---

40 Interpretive use also includes the use of sentences to "represent an assumption, without attributing this assumption to anyone" (Sperber & Wilson, 1986b, p. 229).

theoretic account also predicts, it can be uttered ironically in a rather contrived context in which it is manifest that someone has said or thinks something that entails that the car doesn't have broken windows: e.g. in response to someone smugly saying 'There's practically no crime in *this* neighbourhood and certainly no one here would break into a car to steal the radio.' (Sperber & Wilson, 1986b, pp. 240–241; Wilson, 2006, p. 1732)

This theory of irony is logically independent from the core of relevance theory (and was originally proposed, in a slightly different form, before the core: Sperber & Wilson, 1981). One could, therefore, adopt Sperber and Wilson's theory of irony without relevance theory's core assumptions[41]. Conversely, the relevance theoretic account of irony could be abandoned without giving up any of the core assumptions of relevance theory, as (of course) it might be if it comes into conflict with observation[42].

Making the descriptive/interpretive distinction do explanatory work has been fruitful in the development of a number of other auxiliary hypotheses in relevance theory. In early relevance theory, loose use, hyperbole and metaphor were understood as a kind of interpretive use in which the speaker's utterance is an approximate interpretation of her own thought[43] (Sperber & Wilson, 1986b, p. 231ff; Sperber & Wilson, 1986a), although that account has now been abandoned in favour of one in terms of *ad hoc* concepts (discussed below). Papafragou attempts a relevance theoretic account of metonymy as a type of interpretive use (1996). Finally, combined with a distinction between the desirable and the actual, the descriptive/interpretive distinction underlies the relevance theoretic account of mood and non-declarative sentences: in assertion there is a descriptive relation between speaker's thought and world; in imperatives the speaker's thought describes a desirable state of affairs; and in interrogatives the speaker's thought is in an interpretive relation to desirable thoughts (Sperber & Wilson, 1986b, pp. 231, 243–254; Wilson & Sperber, 1988).

4.5 Ad hoc concepts

The use of the notion of *ad hoc* concepts is a recent development. Relevance theory postulates that strict and literal utterances, loose use, hyperbole and metaphor are not qualitatively distinct phenomena but belong to a continuum. As noted, this was originally explained in terms of interpretive use (Sperber & Wilson, 1986a), but Carston (1997a), and Sperber and Wilson (1998b) now advocate an account in terms of *ad hoc* concepts.

Consider the utterance in (21) made in a context in which it is clear that Mary is talking about her husband, who is not canonised. What she expresses is not that Peter is a strict and literal saint, but something else: perhaps that he is very considerate and self-sacrificing.

(21) Mary: Peter is a saint.

Embedding under logical operators suggests that this is a matter of the proposition expressed, rather than an implicature. A utterance of (22) as a response to (21) is not a denial that Peter is a strict-and-literal saint, but that he is a very nice, considerate etc. individual.

---

41 On the other hand, the classical account of irony is incompatible with the communicative principle. A speaker cannot generally communicate just the opposite of what her words mean because it would cause the hearer gratuitous effort, given that she could just have said what she meant.

42 In fact, it accounts well for intuitions about the examples in the literature (Sperber & Wilson, 1998a; Wilson, 2006), and has received some corroboration from developmental evidence (Happé, 1993).

43 All utterances are interpretive in this way, since an utterance is meant to represent a thought of the speaker's. What relevance theory usually calls interpretive use is use in which that thought is itself interpretive: i.e. resembles another thought, or an utterance, rather than a proposition which describes a state of affairs (Sperber & Wilson, 1986b, p. 231.)

(22) Peter's no saint. He always does what he prefers and makes it look like a huge sacrifice.

Therefore relevance theory postulates that the proposition expressed by (21) is PETER IS A SAINT*, where SAINT* is an *ad hoc* concept: a distinct concept from the lexically encoded concept SAINT, and accessed/constructed by pragmatic inference sensitive to the specific occasion. *Ad hoc* concepts may be broader or narrower than the lexicalised concepts from which they derive. In examples such as (23) and J.L. Austin's (24), the concepts communicated (MILES* and HEXAGONAL*) apply to broader sets than the lexicalised concepts: e.g. France is not HEXAGONAL, but it is HEXAGONAL*, along with many other objects that are not strictly speaking six-sided, but are close enough.

(23) It's miles to the canteen!
(24) France is hexagonal.

Combined broadening and narrowing is seen in (21). The concept SAINT* denotes a set that is both broader than the lexicalised concept (it includes individuals such as Peter, who are not literally saints) and narrower (it will exclude any literal, canonised saints who were not considerate, self-sacrificing etc.).

Recourse to the notion of *ad hoc* concepts has become the primary strategy in relevance theoretic lexical pragmatics. Relevance theorists now argue that lexical pragmatic adjustment is nearly ubiquitous, and "fine-tunes the interpretation of virtually every word." (Carston & Powell, 2006, p. 345)

5 Concluding remarks

Both the core and auxiliary assumptions of relevance theory have developed during its history[44]. To conclude I briefly set out two important early developments in the core, the first of which was mainly driven by the desire to maximise the simplicity, coherence and symmetry of the theory, the second by the aim of bringing relevance theory into line with a development in a related field[45].

The presumption of optimal relevance given in section 3.6 above is stronger than the one originally put forward (Sperber & Wilson, 1986b, p. 158). Sperber and Wilson present and argue for the updated formulation in the postface added to the second edition of the 'Relevance' (1995, p. 267ff.). The original formulation was not symmetrical in effort and effects: its clause b is the presumption that the speaker will maximise the relevance of the utterance, but it treats the intended interpretation (and therefore cognitive effects) as given, so this amounts to an expectation that effort will be minimised. The revision, then, is largely motivated by considerations of simplicity and generality (Sperber & Wilson, 1995, p. 270); although Sperber and Wilson also argue that it increases the predictive power of relevance theory (1995, p. 270).

Another important development in the core of relevance theory is the move to the view that there is a dedicated inferential mechanism for utterance interpretation. In early work, Sperber and Wilson say on the one hand that utterance interpretation appears to be "an ordinary central thought process… relatively unspecialised" (1986b, p. 116) and non-modular (1986b, p. 69), but, on the other hand, suggest that it is carried out by "a less-than-perfect heuristic" (1986b, p. 45) one among "a number of heuristics, some of them innate, others developed through experience, aimed at picking out relevant phenomena"

---

44 For developments in relevance theory see Sperber and Wilson (1995), Carston and Powell (2006) and Clark (2011).

45 Lakatos is sometimes interpreted as saying that the core of a research programme should never change (e.g. Godfrey-Smith, 2003, p. 105). However, his 'negative heuristic' (see note 5 above) merely forbids changing the core in response to empirical problems. It is compatible with this that there be changes to the core not motivated by a direct clash with empirical evidence. Indeed, if changes to the core turn out to be mainly motivated by other considerations that would provide some corroboration for his view.

(1987, p. 703) and argue that analogies with the slow, deliberative reasoning involved in scientific theorising are unhelpful (1986b, p. 117). In more recent work, they argue that there is a mental module dedicated to utterance interpretation (see also Carston, 1997b; 2002). This change reflects considerable rethinking within psychology: of both the nature of central cognition and of the concept of a mental module, and in particular, Sperber's proposal of the massive modularity thesis (Sperber, 1994a). It is also partly prompted by the emergence of the view that human beings have dedicated 'theory of mind' or 'mindreading' abilities (Wimmer & Perner, 1983; Baron-Cohen, Leslie & Frith, 1985; Wellman, 1990): Sperber and Wilson (2002) argue that the comprehension module is related to, but distinct from, the general mindreading module. It should be no surprise that significant developments in cognitive science are reflected in changes in the core of relevance theory, given that the main purpose of relevance theory is to provide a psychologically realistic account of communication.

References

Allott, N. 2006. Game theory and communication. In A. Benz, G. Jäger, & R. van Rooij (Eds.), *Game Theory and Pragmatics*. (pp. 123–151). Basingstoke: Palgrave Macmillan.

Allott, N. 2008. *Pragmatics and Rationality.* Unpublished PhD thesis, University of London.

Baron-Cohen, S., A. M. Leslie, & U. Frith. 1985. Does the autistic child have a "theory of mind"? *Cognition*, *21*(1), 37–46.

Barrett, H. C. 2005. Enzymatic computation and cognitive modularity. *Mind & Language*, *20*(3), 259–287.

Bezuidenhout, A. 1997. Pragmatics, semantic undetermination and the referential/attributive distinction. *Mind*, *106*(423), 375.

Bezuidenhout, A., & M. S. Sroda. 1998. Children's use of contextual cues to resolve referential ambiguity: An application of Relevance Theory. *Pragmatics & Cognition*, *6*(1-2), 265–299.

Bird, A., & E. Tobin. 2010. Natural kinds. In N. Z. Edward (Ed.), *The Stanford Encyclopedia of Philosophy*. (Summer 2010 ed.).

Blakemore, D. 1987. *Semantic Constraints on Relevance.* Oxford: Basil Blackwell.

Blakemore, D. 2000. Indicators and procedures: *nevertheless* and *but. Journal of Linguistics*, *36*(03), 463–486.

Blakemore, D. 2002. *Relevance and Linguistic Meaning: The Semantics and Pragmatics of Discourse Markers.* Cambridge: Cambridge University Press.

Blakemore, D. 2004. Discourse markers. In L. R. Horn, & G. L. Ward (Eds.), *The Handbook of Pragmatics*. (pp. 221–240). Malden, Mass: Blackwell.

Blakemore, D., & R. Carston. 2005. The pragmatics of sentential coordination with *and. Lingua*, *115*(4), 569-589.

Braine, M. D. S. 1978. On the relation between the natural logic of reasoning and standard logic. *Psychological Review*, *85*(1), 1–21.

Braine, M. D. S., & D. P. O'Brien (Eds.). 1998. *Mental Logic.* Mahwah, N.J: L. Erlbaum Associates.

Breheny, R., N. Katsos, & J. Williams. 2006. Are generalised scalar implicatures generated by default? An on-line investigation into the role of context in generating pragmatic inferences. *Cognition*, *100*(3), 434-463.

Carston, R. 1988. Implicature, explicature and truth-theoretic semantics. In R. Kempson (Ed.), *Mental Representations: The Interface Between Language and Reality*. (pp. 155–181). Cambridge: Cambridge University Press.

Carston, R. 1993. Conjunction, explanation and relevance. *Lingua*, *90*, 27–48.

Carston, R. 1997a. Enrichment and loosening: Complementary processes in deriving the proposition expressed. *Linguistische Berichte*, *8*, 103–127.

Carston, R. 1997b. Relevance-theoretic pragmatics and modularity. *UCL Working Papers in Linguistics*, *9*, 29–53.

Carston, R. 1998. Informativeness, relevance and scalar implicature. In R. Carston, & S. Uchida (Eds.), *Relevance theory: applications and implications*. (pp. 179–236). Amsterdam: J. Benjamins.

Carston, R. 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication.* Oxford: Blackwell.

Carston, R. 2010a. Explicit communication and 'free' pragmatic enrichment. In B. Soria, & E. Romero (Eds.), *Explicit Communication: Robyn Carston's Pragmatics*. (pp. 217–285). Basingstoke: Palgrave Macmillan.

Carston, R. 2010b. XIII-Metaphor: *Ad hoc* concepts, literal meaning and mental images. *Proceedings of the Aristotelian Society (Hardback)*, *110*(3pt3), 295-321.

Carston, R., & G. Powell. 2006. Relevance theory: New directions and developments. In E. LePore, & B. C. Smith (Eds.), *The Oxford Handbook of the Philosophy of Language*. (pp. 341–360). Oxford ; New York: Oxford University Press.

Cherniak, C. 1981. Minimal rationality. *Mind*, *90*(358), 161–183.

Chevallier, C., I. Noveck, F. Happé, & D. Wilson. 2011. What's in a voice? Prosody as a test case for the Theory of Mind account of autism. *Neuropsychologia*, *49*(3), 507–517.

Chevallier, C., D. Wilson, F. Happé, & I. Noveck. 2010. Scalar inferences in Autism Spectrum Disorders. *Journal of Autism and Developmental Disorders*, *40*(9), 1104–1117.

Clark, B. 1996. Stylistic analysis and relevance theory. *Language and Literature*, *5*(3), 163–178.

Clark, B. 2011. Recent developments in relevance theory. In D. Archer, & P. Grundy (Eds.), *The Pragmatics Reader*. (pp. 129–137). Oxford: Routledge.

Cohen, L. J. 1971. Some remarks on Grice's views about the logical particles of natural language. In Y. Bar-Hillel (Ed.), *Pragmatics of Natural Languages*. (pp. 50–68). New York: Humanities Press.

Dobzhansky, T. 1964. Biology, molecular and organismic. *American Zoologist*, *4*(4), 443–452.

Emlen, J. M. 1966. The role of time and energy in food preference. *The American Naturalist*, *100*(916), 611–617.

Escandell-Vidal, V., M. Leonetti, & A. Ahern (Eds.). 2011. *Procedural Meaning: Problems and Perspectives*. Bingley: Emerald.

Evans, J. S. B. T., & D. E. Over. 1996. *Rationality and Reasoning*. Hove: Psychology Press.

Fodor, J. A. 1975. *The Language of Thought*. New York: Crowell.

Fodor, J. A. 1985. Fodor's guide to mental representation: The intelligent auntie's vade-mecum. *Mind*, *94*(373), 76–100.

Forster, K. I., & S. M. Chambers. 1973. Lexical access and naming time. *Journal of Verbal Learning and Verbal Behavior*, *12*(6), 627–635.

Gigerenzer, G. 2004. Striking a blow for sanity in theories of rationality. In M. Augier, & J. G. March (Eds.), *Models of a Man: Essays in Memory of Herbert A. Simon*. (pp. 389–409). Cambridge, Mass: MIT Press.

Gigerenzer, G., & D. G. Goldstein. 1996. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650–669.

Gigerenzer, G., & P. M. Todd. 1999. *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Godfrey-Smith, P. 2003. *Theory and Reality: An Introduction to the Philosophy of Science*. Chicago: University of Chicago Press.

Grice, P. 1967. Logic and conversation: William James lectures. *Studies in the Way of Words (1989)*. (pp. 1–143). Cambridge, Mass: Harvard University Press.

Grice, P. 1971. Intention and Uncertainty. *Proceedings of the British Academy*, *LVIII*, 263–279.

Grice, P. 1975. Logic and conversation. In P. Cole, & J. Morgan (Eds.), *Syntax & Semantics 3: Speech Acts*. (pp. 41–58). New York: Academic Press.

Grice, P. 1989. *Studies in the Way of Words*. Cambridge, Mass: Harvard University Press.

Gutt, E.-A. 1991. *Translation and Relevance: Cognition and Context*. Oxford: Blackwell.

Hacking, I. 1979. Imre Lakatos's philosophy of science. *The British Journal for the Philosophy of Science*, *30*(4), 381–402.

Happé, F. 1993. Communicative competence and theory of mind in autism: A test of relevance theory.

Cognition, *48*(2), 101–119.

Happé, F., & E. Loth. 2002. Theory of mind and tracking speakers intentions. *Mind & Language*, *17*, 24–36.

Ifantidou, E. 2001. *Evidentials and Relevance.* Amsterdam: J. Benjamins.

Iten, C. 2005. *Linguistic Meaning, Truth Conditions and Relevance: The Case of Concessives.* New York: Palgrave Macmillan.

Jary, M. 2007. Are explicit performatives assertions? *Linguistics and Philosophy*, *30*(2), 207-234.

Jary, M. 2010. *Assertion.* Basingstoke: Palgrave Macmillan.

Johnson-Laird, P. N. 1983. *Mental Models.* New York: Cambridge University Press.

Jorgensen, J., G. A. Miller, & D. Sperber. 1984. Test of the mention theory of irony. *Journal of Experimental Psychology: General*, *113*(1), 112–120.

Kempson, R. M. 1975. *Presupposition and the Delimitation of Semantics.* Cambridge: Cambridge University Press.

Kjøll, G. 2010. Content similarity and communicative success. *International Review of Pragmatics*, *2*(1), 21–45.

Lakatos, I. 1968. Criticism and the methodology of scientific research programmes. *Proceedings of the Aristotelian Society New Series*, *69*, 149–186.

Lakatos, I. 1970. Falsification and the methodology of scientific research programmes. In I. Lakatos, & A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*. (pp. 91–195). Cambridge: Cambridge University Press.

Lakatos, I., J. Worrall, & G. Currie. 1978. *Philosophical Papers.* Cambridge: Cambridge University Press.

Lavie, N. 1995. Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(3), 451–468.

Lavie, N. 2001. The role of capacity limits in selective attention: Behavioural evidence and implications for neural activity. In J. Braun, & C. Koch (Eds.), *Visual attention and cortical circuits*. (pp. 49–68). Cambridge, Mass.: MIT Press.

Levinson, S. C. 2006. On the human 'interactional engine'. In N. J. Enfield, & S. C. Levinson (Eds.), *Roots of Human Sociality: Culture, Cognition and Interaction*. (pp. 39–69). Oxford: Berg.

Matsui, T. 2000. *Bridging and Relevance.* Amsterdam: J. Benjamins.

Meyer, D. E., & R. W. Schvaneveldt. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*(2), 227–234.

Moura, A. C. d. A., & P. C. Lee. 2004. Capuchin stone tool use in Caatinga dry forest. *Science*, *306*(5703), 1909–1909.

Neale, S. 1992. Paul Grice and the philosophy of language. *Linguistics and Philosophy*, *15*(5), 509–559.

Neely, J. H. 1991. Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner, & G. W. Humphreys (Eds.), *Basic Processes in Reading : Visual Word Recognition*. (pp. 264–336). Hillsdale, N.J.: L. Erlbaum Associates.

Newell, A., & H. A. Simon. 1976. Computer science as empirical inquiry: Symbols and search. *Communications of the Association for Computing Machinery*, *19*(3), 113–126.

Nicolle, S., & B. Clark. 1999. Experimental pragmatics and what is said: A response to Gibbs and Moise. *Cognition*, *69*(3), 337–354.

Noveck, I. A., & D. Sperber (Eds.). 2004. *Experimental Pragmatics*. New York: Palgrave Macmillan.

Noveck, I. A., & D. Sperber. 2007. The why and how of experimental pragmatics: The case of 'scalar inferences'. In N. Burton-Roberts (Ed.), *Pragmatics*. (pp. 184–212). Basingstoke: Palgrave

Macmillan.

Papafragou, A. 1996. On metonymy. *Lingua*, *99*(4), 169–195.

Pashler, H. E. 1998. *The Psychology of Attention.* Cambridge, Mass: MIT Press.

Pilkington, A. 2000. *Poetic Effects : A Relevance Theory Perspective.* Amsterdam ; Philadelphia: J. Benjamins Pub.

Posner, R. 1980. Semantics and pragmatics of sentence connectives in natural language. In J. R. Searle, F. Kiefer, & M. Bierwisch(pp. 168–203). Dordrecht, Boston, London: Riedel.

Powell, G. 2001. The referential-attributive distinction: A cognitive account. *Pragmatics and Cognition*, *9*(1), 69-98.

Powell, G. 2010. *Language, Thought and Reference.* Basingstoke: Palgrave Macmillan.

Quine, W. V. 1969. Natural kinds. *Ontological Relativity, and Other Essays.* (pp. 114–138). Princeton, N.J.: Columbia University Press.

Rips, L. J. 1983. Cognitive processes in propositional reasoning. *Psychological Review*, *90*(1), 38–71.

Rouchota, V. 1992. On the referential/attributive distinction. *Lingua*, *87*, 137-167.

Selfridge, O. G., & U. Neisser. 1960. Pattern recognition by machine. *Scientific American*, *203*(2), 60–68.

Simon, H. A. 1957. *Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in a Social Setting.* New York: Wiley.

Simon, H. A. 1990. Invariants of human behavior. *Annual Review of Psychology*, *41*(1), 1–19.

Sloman, S. A. 1996. The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*(1), 3–22.

Sperber, D. 1994a. The modularity of thought and the epidemiology of representations. In L. A. Hirschfeld, & S. A. Gelman (Eds.), *Mapping the Mind: Domain Specificity in Cognition and Culture.* (pp. 39–67). Cambridge: CUP.

Sperber, D. 1994b. Understanding verbal understanding. In J. Khalfa (Ed.), *What is Intelligence?* (pp. 179–198). Cambridge: Cambridge University Press.

Sperber, D. 2005. Modularity and Relevance: How Can a Massively Modular Mind Be Flexible and Context-Sensitive? In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The Innate Mind: Structure and Contents.* (pp. 53–68). Oxford University Press.

Sperber, D., F. Cara, & V. Girotto. 1995. Relevance theory explains the selection task. *Cognition*, *57*(1), 31–95.

Sperber, D., F. Clément, C. Heintz, O. Mascaro, H. Mercier, G. Origgi, & D. Wilson. 2010. Epistemic vigilance. *Mind & Language*, *25*(4), 359–393.

Sperber, D., & D. Wilson. 1981. Irony and the use-mention distinction. In P. Cole (Ed.), *Radical Pragmatics.* New York: Academic Press.

Sperber, D., & D. Wilson. 1986a. Loose talk. *Proceedings of the Aristotelian Society New Series*, *86*, 153–171.

Sperber, D., & D. Wilson. 1986b. *Relevance: Communication and Cognition (2nd Ed. 1995).* Oxford: Blackwell.

Sperber, D., & D. Wilson. 1987. Précis of 'Relevance: Communication and Cognition'. *Behavioral and Brain Sciences*, *10*, 697–754.

Sperber, D., & D. Wilson. 1995. Postface. *Relevance: Communication and Cognition.* (2nd ed., pp. 255–279). Oxford: Blackwell.

Sperber, D., & D. Wilson. 1996. Fodor's frame problem and relevance theory (reply to Chiappe & Kukla). *Behavioral and Brain Sciences*, *19*(3), 530–532.

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*. (pp. 283–293). Amsterdam: J. 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*. (pp. 184–200). Cambridge: Cambridge University Press.

Sperber, D., & D. Wilson. 2002. Pragmatics, modularity and mind-reading. *Mind and Language*, *17*(1–2), 3–23.

Sperber, D., & D. Wilson. 2008. A deflationary account of metaphors. In R. W. Gibbs (Ed.), *The Cambridge Handbook of Metaphor and Thought*. (pp. 84–108). New York: Cambridge University Press.

Stephens, D. W., J. S. Brown, & R. C. Ydenberg (Eds.). 2007. *Foraging: Behavior and Ecology*. Chicago, Ill.: University of Chicago Press.

Stephens, D. W., & J. R. Krebs. 1986. *Foraging Theory*. Princeton, N.J.: Princeton University Press.

Todd, P. M., & G. Gigerenzer. 2000. Précis of 'Simple heuristics that make us smart'. *Behavioral & Brain Sciences*, *23*(5), 727–41; discussion 742–80.

Unger, C. 2006. *Genre, Relevance, and Global Coherence : The Pragmatics of Discourse Type*. New York: Palgrave Macmillan.

van der Henst, J.-B., L. Carles, & D. Sperber. 2002. Truthfulness and relevance in telling the time. *Mind and Language*, *17*(5), 457–466.

van der Henst, J.-B., G. Politzer, & D. Sperber. 2002. When is a conclusion worth deriving? A relevance-based analysis of indeterminate relational problems. *Thinking & Reasoning*, *8*(1), 1–20.

Vega, M., R. E. 2007. *Creativity and Convention: The Pragmatics of Everyday Figurative Speech*. Amsterdam: John Benjamins.

Wedgewood, D. 2007. Shared assumptions: semantic minimalism and relevance theory. *Journal of Linguistics*, *43*, 647–681.

Wellman, H. M. 1990. *The Child's Theory of Mind*. Cambridge, Mass: MIT Press.

Wharton, T. 2009. *Pragmatics and Non-Verbal Communication*. Cambridge: Cambridge University Press.

Wilson, D. 1975a. Presupposition, assertion, and lexical items. *Linguistic Inquiry,* 95–114.

Wilson, D. 1975b. *Presuppositions and Non-Truth-Conditional Semantics*. London: Academic Press.

Wilson, D. 2006. The pragmatics of verbal irony: Echo or pretence? *Lingua*, *116*(10), 1722–1743.

Wilson, D. 2009. Relevance theory. In L. Cummings (Ed.), *The Pragmatics Encyclopedia*. (pp. 393–399). London: Routledge.

Wilson, D. 2011. The conceptual-procedural distinction: Past, present and future. In V. Escandell-Vidal, M. Leonetti, & A. Ahern (Eds.), *Procedural Meaning: Problems and Perspectives*. (pp. 5–31). Bingley: Emerald.

Wilson, D., & T. Matsui. 1998. Recent approaches to bridging: Truth, coherence, relevance. *UCL Working Papers in Linguistics*, *10*, 173–200.

Wilson, D., & D. Sperber. 2012. *Meaning and Relevance*. Cambridge: Cambridge University Press.

Wilson, D., & D. Sperber. 1981. On Grice's theory of conversation. In P. Werth (Ed.), *Conversation and Discourse*. (pp. 155–178). London: Croom Helm.

Wilson, D., & D. Sperber. 1988. Mood and the analysis of non-declarative sentences. In J. Dancy, J. M. E. Moravcsik, & C. C. W. Taylor (Eds.), *Human Agency: Language, Duty, and Value; Philosophical Essays in Honor of J. O. Urmson*. (pp. 77–101). Stanford, Calif: Stanford University Press.

Wilson, D., & D. Sperber. 1992. On verbal irony. *Lingua*, *87*(1/2), 53–76.

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

Wilson, D., & D. Sperber. 2002. Truthfulness and relevance. *Mind*, *111*(443), 583–632.

Wilson, D., & D. Sperber. 2004. Relevance theory. In L. R. Horn, & G. L. Ward (Eds.), *The Handbook of Pragmatics*. (pp. 607–632). Malden, Mass: Blackwell.

Wimmer, H., & J. Perner. 1983. Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition*, *13*(1), 103–128.