

Testing the cognitive and communicative principles of relevance

Jean-Baptiste Van Der Henst, Dan Sperber

► **To cite this version:**

Jean-Baptiste Van Der Henst, Dan Sperber. Testing the cognitive and communicative principles of relevance. Ira Noveck, Dan Sperber. Experimental Pragmatics, Palgrave, 2004. <ijn_00000452>

HAL Id: ijn_00000452

https://jeannicod.ccsd.cnrs.fr/ijn_00000452

Submitted on 5 Apr 2004

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

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

To appear in Noveck, I. & Sperber, D. (eds) *Experimental Pragmatics*, Palgrave

Testing the cognitive and communicative principles of relevance

Jean Baptiste Van der Henst & Dan Sperber

1. Introduction

A general theory is testable not directly but through consequences it implies when it is taken together with auxiliary hypotheses. The test can be weaker or stronger depending, in particular, on the extent to which the consequences tested are specifically entailed by the theory (as opposed to being mostly entailed by the auxiliary hypotheses and being equally compatible with other general theories). The earliest experimental work based on relevance theory (Jorgensen, Miller and Sperber 1984, Happé 1993) tested and confirmed Sperber & Wilson (1981) echoic account of irony (and much experimental work done since on irony has broadly confirmed it and refined it further). While this account of irony is part and parcel of relevance theory, it is nevertheless compatible with different pragmatic approaches. The experimental confirmation of this account, therefore, provides only weak support for relevance theory as a whole. More recent experimental work has made explicit, tested and confirmed other and more specific and central consequences of relevance theory (e.g. Sperber, Cara & Girotto, 1995; Politzer, 1996; Gibbs & Moise, 1997; Hardman 1998; Nicolle & Clark 1999; Matsui, 2000, 2001; Girotto, Kimmelmeir, Sperber & Van der Henst, 2001; Noveck, 2001; Noveck, Bianco & Castry, 2001; Van der Henst, Politzer and Sperber, 2002, Van der Henst, Carles & Sperber, 2002, Noveck & Posada, 2003; Ryder & Leinonen, 2003).

Here we review experiments that test consequences of the most central tenets of the theory, namely the cognitive and the communicative principle of relevance.

2. The basic tenets of relevance theory

Relevance, as characterised in relevance theory, is a property of inputs to cognitive processes. These inputs include external stimuli (for instance utterances) and internal representations (for instance memories or conclusions from inferences that may then be used as premises for further inferences). When is an input relevant? An input is relevant to an individual when processing it in a context of previously available assumptions yields positive cognitive effects, that is, improvements to the individual's knowledge that could not be achieved from processing either the context on its own, or the new input on its own. These improvements may consist in the derivation of contextual implications, in the confirmation of uncertain assumptions, in the correction of errors, and also, arguably, in the reorganisation of knowledge so as to make it more appropriate for future use.

Inputs are not just relevant or irrelevant; when relevant, they are more or less so. A relatively high degree of relevance is what makes some inputs worth processing. Many of the potential inputs competing for an individual's processing resources at a given time may offer a modicum of relevance, but few are likely to be relevant enough to deserve attention. What makes these worth processing is, to begin with, that they yield comparatively higher cognitive effects. However, two inputs yielding the same amount of cognitive effect may differ in the amount of processing effort¹ required to produce

¹ "Effort" as used here refers here to any expenditure of energy in the pursuit of a goal. It is not restricted to conscious effort.

this effect. Obviously, the lesser the effort, the better. If relevance is what makes an input worth processing, then the relevance of an input is not just a matter of the cognitive effect it yields but also of the mental effort it requires. Hence the characterisation of relevance in terms of *effect* and *effort*:

- (1) *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.

Here is a simplified illustration of how the relevance of alternative inputs might be compared in terms of effort and effect. Suppose you want to take the next train to Bordeaux and compare statements (2)-(4) (assumed to be uttered by a reliable informer):

- (2) The next train to Bordeaux is at 3:24 pm
- (3) The next train to Bordeaux is after 3pm
- (4) The next train to Bordeaux is 36 minutes before 4pm

All three statements would be relevant to you, but (2) would be more relevant than either (3) or (4). Statement (2) would be more relevant than (3) for reasons of cognitive effect: (2) entails (3), and therefore yields all the conclusions derivable from (2), and more besides, and these extra conclusions themselves have practical consequences for the

planning of your trip. Statement (2) would be more relevant than (4) for reasons of processing effort: although (2) and (4) are logically equivalent, and therefore yield exactly the same cognitive effects, these effects are easier to derive from (2) than from (4), which requires an additional effort of calculus with no additional benefit whatsoever (in the ordinary situation envisaged). More generally, when similar amounts of effort are required by two alternative inputs, the effect factor is decisive in determining degrees of relevance, and when similar amounts of effect are achievable, the effort factor is decisive. In experimental work, as we will illustrate, this makes it relatively easy to manipulate the relevance of stimuli across conditions by keeping the effort factor constant and modifying the effect factor or, conversely, by keeping the effect factor constant and modifying the effort factor.

Relevance theory claims that, because of the way their cognitive system has evolved, humans have an automatic tendency to maximise relevance. As a result of constant selection pressure towards efficiency, perceptual mechanisms tend automatically to pick out potentially relevant stimuli, memory mechanisms tend automatically to store and, when appropriate, retrieve potentially relevant pieces of knowledge, and inferential mechanisms tend spontaneously to process these inputs in the most productive way. This universal tendency is described in the First, or Cognitive, Principle of Relevance:

(5) *Cognitive Principle of Relevance*

Human cognition tends to be geared to the maximisation of relevance.

This spontaneous tendency to maximise relevance makes it possible to predict to some extent to which available stimuli people will pay attention and how they will process them.

There is a wealth of evidence in the experimental study of attention and memory that could be re-analysed in order to see to what extent it supports the cognitive principle of relevance. This is not our field of expertise, but the challenge there, we surmise, would be not so much to find support as to find support that is specific enough to relevance theory, in other words to find predictions that follow from the cognitive principle of relevance but not—or not as directly—from standard psychological approaches to attention and memory. In other areas, the study of inference and that of communication in particular, the cognitive principle does have consequences that are far from trivial. Some of these consequences in the domain of category-based induction have been explored by Medin et al (in press). In section 2 of this chapter we will present experimental tests of consequences based on work by Van der Henst and his collaborators on relational reasoning.

Relevance theory has been mostly an exploration of the implications of the second, Communicative Principle of Relevance for human verbal communication. The human tendency to maximise relevance makes it possible not only to predict some of other people's cognitive processes, but also to try to influence them—how indeed could you aim at influencing people if you had no way to predict how your behaviour would affect their thought? Human intentional communication, and in particular verbal communication, involves the attribution, by the communicator and the addressee, of mental states to one another. This attribution is greatly helped by the relative predictability of relevance-guided cognitive processes. In particular a speaker must

intend and expect that the hearer will pay attention to the utterance produced. If attention tends automatically to go to inputs that seem relevant enough to be worth processing, then it follows that, to succeed, the speaker must intend and expect her utterance to be seen as relevant enough by the hearer she is addressing. By the very act of speaking to him, the communicator therefore encourages the hearer to presume that the utterance is so relevant. This is the basis for the Communicative Principle of Relevance:

(6) *Communicative Principle of Relevance*

Every utterance conveys a presumption of its own optimal relevance.

An utterance, so the theory claims, conveys not just a vague expectation, but a precise presumption of relevance, which the notion of “optimal relevance” captures:

(7) *Optimal relevance*

An utterance is optimally relevant to the hearer just in case:

- (a) It is relevant enough to be worth the hearer’s processing effort;
- (b) It is the most relevant one compatible with the speaker’s abilities and preferences.

According to clause (7a) of this definition, the hearer is entitled to expect the utterance to be at least relevant enough to be worth processing, which means (given the cognitive principle of relevance) that the utterance should be more relevant than any alternative input available at the time.

Is the hearer entitled to higher expectation than this (already high) minimum level spelled out in clause (7a)? The speaker wants to be understood. It is therefore in her interest to make her utterance as easy as possible to understand, and to provide evidence not just for the cognitive effects she aims to achieve in the hearer but also for further cognitive effects which, by holding his attention, will help her achieve her goal. Speakers, however, are not omniscient, and they cannot be expected to go against their own interests and preferences in producing an utterance. There may be relevant information that they are unable or unwilling to provide, and wordings that would convey their meaning more economically, but that they are unable to think of at the time, or are unwilling to use (for reason of propriety for instance). All this is spelled out in clause (7b) of the definition of optimal relevance, which states that the ostensive stimulus is the most relevant one (i.e. yielding the greatest effects, in return for the smallest processing effort) that the communicator is *able and willing* to produce.

The Communicative Principle of Relevance justifies a specific inferential procedure for interpreting an utterance, that is for discovering what the speaker meant by uttering it:

- (8) *Relevance-guided comprehension procedure*
- a. Follow a path of least effort in constructing and testing interpretive hypotheses (regarding disambiguation, reference resolutions, implicatures, etc.).
 - b. Stop when your expectations of relevance are satisfied.

Given clause (7b) of the definition of optimal relevance, it is reasonable for the hearer to follow a path of least effort because the speaker is expected (within the limits of her abilities and preferences) to make her utterance as easy as possible to understand. Since relevance varies inversely with effort, the very fact that an interpretation is easily accessible gives it an initial degree of plausibility. It is also reasonable for the hearer to stop at the first interpretation that satisfies his expectations of relevance, because there should never be more than one. A speaker who wants her utterance to be as easy as possible to understand should formulate it (within the limits of her abilities and preferences) so that the first interpretation to satisfy the hearer's expectation of relevance is the one she intended to convey. An utterance with two apparently satisfactory competing interpretations would cause the hearer the unnecessary extra effort of choosing between them, and, because of this extra effort, the resulting interpretation (if there were one) could never satisfy clause (7b) of the definition of optimal relevance. Thus, when a hearer following the path of least effort arrives at an interpretation that satisfies his expectations of relevance, he should, in the absence of contrary evidence, adopt it. Since comprehension is a non-demonstrative inference process, this interpretation of the speaker's meaning may be erroneous. Still, it is the most plausible interpretation in the circumstances.

The hypothesis that hearers spontaneously follow the relevance-guided comprehension procedure spelled out in (8) can be experimentally tested by manipulating either the effort factor and, in particular, by changing the order of accessibility of various interpretations. It can also be tested by manipulating the effect factor and thereby making specific interpretation more or less likely to satisfy the hearer's expectations of relevance. This, as we will illustrate in section 3, is what

Giroto, Sperber and their collaborators have done in a series of experiments with the Wason Selection Task.

Most work in relevance theory so far has been focused on utterance interpretation rather than on utterance production. The theory, however, has testable implications regarding the production process. Speakers often fail to be relevant to their audience, and sometimes do not even make the effort to be relevant. Still, utterances couldn't effectively convey the presumption of their own relevance unless speakers were, most of the time, aiming at optimal relevance and achieving it often enough. In section 4, we describe a series of experiments that were aimed at testing to what extent speakers were actually aiming at optimal relevance.

3. Testing the Cognitive Principle of Relevance with relational reasoning tasks

In most studies on reasoning, psychologists analyse participants' successful or unsuccessful performance in reasoning tasks. They look at the percentages of correct conclusions or at the time taken to draw such a conclusion. They investigate factors that impede or enhance correct performance, such as the premises' content, the premises' complexity, task instructions, or IQ. They use this evidence to test various theories of the inferential machinery that underlies our reasoning ability. Some argue that people reason by constructing mental models of the premises (Johnson-Laird & Byrne, 1991). Others support the idea that people reason by applying general inference rules (Rips, 1994; Braine & O'Brien, 1998). Yet others have proposed that reasoning relies on domain-specific procedures (Cheng & Holyoak, 1985; Cosmides, 1989).

Relevance theory claims that comprehension is based on a domain-specific inferential procedure, but it is not, in and by itself, a theory of human reasoning. It is, in

fact, compatible with the view that an important role is played in reasoning by mental models, or by inference rules, or by both, or by yet other kinds of procedures in a domain-general or in a domain-specific way.² Nevertheless, relevance theory may make a direct contribution to the study of reasoning by suggesting testable claims not on the procedures (except in the case of comprehension) but on the goals of reasoning processes.

Standard approaches to the study of reasoning have had little to say on what causes people to engage in reasoning—when they are not, that is, requested to do so by an experimenter—, what expectations they have in doing so, and what kind of conclusions satisfy these expectations bringing the process to a close.³ What guides reasoners to infer a specific conclusion? At first sight, one might argue that people aim at inferring a conclusion that *logically follows* from the premises. However, from any given set of premises, an infinity of conclusions logically follows. Most of these valid conclusions are of no interest at all. For instance, nobody would burden one's mind by inferring from the single premise P the logical conclusion Not (not (not (not P))). Harman has formulated this idea as a *principle of clutter avoidance*: “It is not reasonable or rational to fill your mind with trivial consequences of your beliefs, when you have better things to do with your time, as you often do” (Harman, 1995, p. 186).

It is not sufficient for a conclusion to be logically valid in order to be worth inferring. Some valid conclusions are too trivial ever to be derived, and others may be derived in some circumstances and not in others. From the same set of premises, we

² Sperber however has been defending the view that the human mind is “massively modular” (Sperber 1994), and Sperber and Wilson (2002) have argued that linguistic comprehension is modular.

³ See Johnson-Laird and Byrne (1991:20-22) for a notable exception.

might derive one particular conclusion in one situation, another conclusion in a second situation, or no conclusions at all in a third. In a recent study, we proposed that the conclusions that people are inclined to draw are those, if any, that seem relevant enough in the context (Van der Henst, Sperber & Politzer, 2002). This, of course, is a direct consequence of the cognitive principle of relevance.

In this study, we compared so-called “determinate” and “indeterminate” relational problems such as those:

A determinate problem:

A is taller than B

B is taller than C

An indeterminate Problem:

A is taller than B

A is taller than C

Such relational problems have been empirically investigated in many studies (see Evans, Newstead & Byrne, 1993 for a review). Determinate problems are so-called because the one relation between the three terms A, B, and C which is not explicitly described in the premises, that between A and C, is nevertheless inferable from them: in our example, A is taller than C. Indeterminate problem are so-called because the one relation which is not described in the premises is not inferable from them: in our example, B might be taller than C or C might be taller than B. Hence, nothing follows from the premises about the relation between B and C. The goal of most studies on relational problems has been to describe the way in which the premises are being mentally represented and processed by reasoners. Typically, participants have had to answer a specific question like “What is the relation between A and C?” and the evidence consists in the rate of correct answers. The correct answer for the determinate problem above would be: A is

taller than C. The correct answer for indeterminate problem would be: it is impossible to tell. Indeterminate problems tend to yield a lower rate of correct answers than determinate problems.⁴

In our study, our aim was not to assess and explain the relative difficulty of determinate and indeterminate problems. Instead of asking a question about a specific relation between two terms mentioned in the premises, we just asked *what, if anything, follows from the premises?* We were interested in what causes some participants, particularly with indeterminate problems, to answer *nothing follows*.

Not only is it always possible to infer conclusions from a given set of premises, but what is more, some of these conclusions are quite obvious: for instance, from two premises P and Q, their conjunction P-and-Q trivially follows. So when people answer that nothing follows from a given set of premises, either they just fail to see the obvious, or, we suggest, they mean that nothing *relevant* follows. If so, *nothing follows* answer are evidence of people's intuitions of relevance. In particular, if a problem creates the expectation that the most relevant conclusion to be derived should be of a certain type and, at the same time, does not warrant any conclusion of this particular type, people may be tempted to answer that nothing follows. This, we tried to show, is what happens with indeterminate relational problems.

What conclusion could participants expect to infer from two relational premises in the context of a reasoning task? In determinate and indeterminate relational problems such as the examples above, there are three terms, A, B, and C, one type of asymmetric

⁴ Supporters of mental model theory explain this fact by pointing out that the mental representation of indeterminate problems calls for two mental models (to represent the two possible relations between B and C) as opposed to one model for the determinate problems (Byrne & Johnson-Laird, 1989).

and transitive relation, e.g. *taller than*, and therefore three possible relations of this type, in the pairs A-B, B-C, A-C. Two of these relations are described in the premises. Given the communicative principle of relevance, these relations are presumed to be relevant in the context of the task, and, more specifically, the two relations given in the premises are expected to be relevant in allowing the inference of the third relation. Of course, it could be rightly pointed out that, in these experimental situations, the premises on which participants are asked to reason are arbitrary and without relationship with their real-life concerns. Therefore neither the premises nor the conclusions that can be derived from them have any genuine relevance. Still, we would argue, just as participants reason under the pretence that the premises are true (that, say, the premise “Jim is taller than Paul” is about two actual people), they reason under the pretence that the premises, and the conclusions they are expected to derive from them, might be relevant in some ordinary context of knowledge about the individuals or the entities described in the premises. It is not hard, for instance, to pretend that it might be relevant to know that Jim is taller than Paul and that Paul is taller than Dick, and to assume then that it would be relevant to draw the inference that Jim is taller than Dick.

Participants’ expectations of relevance are easily satisfied in the case of determinate problems but not in the case of indeterminate ones, where the relation that is not specified in the premises cannot be inferred from them. Hence, with indeterminate problems, participants may be tempted to answer that nothing follows. This is indeed what we observed. In our studies, 45% of the participants gave a “nothing follows” response to indeterminate problems, while only 8 % did so with determinate problems. This difference in the rate of “nothing follows” answer between determinate and indeterminate problems is, of course, not surprising. However, it had never been

demonstrated before, and, more importantly, only relevance theory provides a simple and direct explanation of this difference. When participants say that nothing follows, what they mean, we surmise, is not that it is impossible to infer anything at all from the two premises, but that it is impossible to derive a conclusion relevant enough to be worth deriving, namely, a conclusion about the third undescribed relation among the three items mentioned in the premises.

Nevertheless, facing a situation where what would be the most relevant conclusion cannot be inferred, about one half of the participants do offer some positive conclusion. Are they giving up on relevance and aiming just for any logically valid conclusion, or are they still guided by considerations of relevance? As we will show, one can find out by examining the specific conclusions they actually derive.

Consider the determinate conclusion *A is taller than B and C*, or equivalently, *A is the tallest*, derived from our Indeterminate Problem. This conclusion is merely a linguistic integration of the premises. It may seem trivial, especially in the context of a reasoning experiment where, generally, participants are eager to demonstrate their reasoning skills to the experimenter. However, a conclusion such as *A is the tallest* may have some relevance of its own. There are ordinary situations where it would be relevant to know which item in a set is above the others with respect to some given property (e.g. who is the tallest?). Actually, in many situations, knowing which item in a set is above all the others with respect to some comparative property is more relevant than knowing the relative position of two other items in the set that are lower on the comparison scale. For instance, suppose you have the choice among three different cars all of which would satisfy your needs and you just want to buy the cheapest. You will probably be more interested in knowing which is the cheapest of the three than in knowing which is the

cheaper of the other two. Hence, inferring *A is more... than B and C* has some relevance since, assuming a quite ordinary context, it can be a step towards inferring further contextual implications (e.g. about which car to buy).

One might query: How can deductively deriving a conclusion and adding it to, or substituting it for, an initial set of premises yield a more relevant point of departure for further reasoning, given that nothing can be derived from this conclusion that wasn't already derivable from the initial premises? In other words, how can such a conclusion be relevant at all, in a context where the premises from which it is derived are given? The fact that relevance is defined not just in terms of effect but also in terms of processing effort provides a simple answer. A set of premises with some deductively derived conclusion added could not carry more cognitive effects than the initial set and thus be more relevant *on the effect side*, but it can be more relevant *on the effort side* by allowing the same effects to be derived with less effort. The deduction of some specific conclusion from a set of premises may be a preliminary and effort-costly necessary step towards deriving cognitive effects from this set of premises. In that case, the conclusion is as relevant as the premises on the effect side and more relevant than the premises on the effort side.

We frequently encounter information which we think is likely to prove useful in the future. We then retain this information, and often process it in such a way as to optimize its potential usefulness. Suppose, for instance, that you arrive in a holiday resort where you plan to spend a month with your family. You learn that there are three doctors in the resort, Smith, Jones, and Williams. You also learn the following two pieces of information: *{Smith is a better doctor than Jones, Jones is a better doctor than Williams}*. At the time, you don't need a doctor, but you might in the future, and would

then want to visit the best doctor in town. So the information is potentially relevant to you. You might just store the two pieces of information above, but from a cognitive point of view it would be more efficient to draw the conclusion: *Smith is the best doctor* straight away. By drawing this conclusion now, you prepare for future circumstances in which you would need a doctor. By adding this conclusion to the two initial premises, you are left with a set of premises for future inference with a greater expected relevance, since its exploitation will require fewer inferential steps. Moreover, if you expect not to need information about the other two doctors, it may be sufficient to remember just the conclusion *Smith is the best doctor*, replacing the initial two-premise set with the single derived conclusion, thus reducing the memory load.

If what makes a conclusion seem relevant is that it spares effort for the possible derivation of cognitive effects, then it follows that the more effort it spares for such possible derivations, the greater will be its perceived relevance. In our initial study (Van der Henst, Sperber & Politzer, 2002), we manipulated the relevance of a relational conclusion of the form *A is more...than B & C* by formulating the premises so as to make the derivation of such a conclusion more or less effortful. In one type of problem, the derivation of this conclusion was very easy and thus the effort saved for the possible derivation of cognitive effects was quite low, whereas with another type of problem, deriving the conclusion was harder, and thus the effort saved was greater. The problems we used were the following⁵:

⁵ All the experiments reported in this section were carried out in French.

Problem 1:

A is taller than B

A is taller than C

Problem 2:

B is taller than A

C is taller than A

In both problems, the relation between B and C is indeterminate. Still, from either problem one can derive a variety of conclusions. For instance, from Problem 1, one can infer conclusion (9a) and (9b), and from Problem 2, one can infer conclusion (10a) and (10b):

(9) (a) '*A is taller than B and C*'

(b) '*B and C are shorter than A*'

(10) (a) '*A is shorter than B and C*'

(b) '*B and C are taller than A*'

With the usual element of pretence involved in the experimental study of reasoning, such conclusions can be seen as having some relevance in that they may facilitate the derivation of further cognitive effects, given some plausible context.

Deriving the single-subject conclusion (9a) from the premises of Problem 1 hardly involves any inferential effort. Since the grammatical subject (*A*) and the comparative term (*taller than*) are the same in the conclusion and in the premises, it amounts just to

merging the two premises in a single sentence. Deriving the single-subject conclusion (10a) from the premises of Problem 2, on the other hand, involves some genuine inferential effort: the grammatical object in the premises (*A*) has to be put in subject position, and the comparative term (*taller than*) has to be converted into its opposite (*shorter than*). It is rather the double-subject conclusion (10b) that amounts to a mere merging of the premises. If participants just went for the less effort-demanding conclusion, they should choose (9a) and (10b). However, if they are guided by considerations of relevance, they should choose (9a) and (10a).

Conclusions (9a) and (9b) are logically equivalent and therefore, in any context, would yield the same effects, and so are and would conclusions (10a) and (10b). However, in most contexts, deriving these effects by using the single-subject conclusions (9a) and (10a) as premises is likely to cause less effort than doing by using the double-subject conclusions (9b) or (10b) as premises. Why? Because most piece of knowledge transmitted, constructed, and stored in human cognition has as its topic a single entity or a single category rather than a pair of entities or categories (for fairly obvious reasons having to do with cognitive efficiency). One is more likely, for instance, to encounter a contextual conditional premise of the form (11a) than of the form (11b):

- (11) (a) '*If A is taller than B and C, then...*'
 (b) '*If B and C are shorter than A, then...*'

From either of (9a) and (9b) and either of (11a) and (11b) as premises, the same conclusions can be derived, but the derivation will be more direct if the minor premise,

i.e. (9a) or (9b), of this conditional syllogism matches the antecedent of the major premise, i.e. (11a) or (11b). In other words, in most realistic contexts, single-subject conclusions such as (9a) and (10a) are likely to prove more relevant than double-subject conclusions such as (9b) and (10b). We predicted therefore that, in both Problem 1 and Problem 2, participants, guided by considerations of relevance, would derive more single-subject than double-subject conclusions.

There is a further reason, specific to the premises of Problem 2, why (10a) should be perceived as more potentially relevant than (10b). It is that the extra effort involved in deriving (10a) as compared to (10b) is effort expended in the right direction. It can be seen as preparatory for the derivation of cognitive effect. This argument does not apply to (9a) and (9b) in Problem 1. As we mentioned, the derivation of (9a), unlike that of (9b), involves almost no effort. In other terms, the derivation of both (9a) and (10a) are steps in the right direction, but the derivation of (10a) is a much bigger step, and therefore a more useful one. This suggests that Problem 2 should be seen as yielding a relevant enough conclusion more frequently than Problem 1.

For the reasons just developed, we expected that participants who produced a conclusion with Problem 1 and 2 would predominantly produce a single-subject conclusion and that there would be more such conclusions, and fewer “nothing follows,” with Problem 2 than with Problem 1. Note that there is nothing intuitively obvious about these predictions, which follow quite directly from the cognitive principle of relevance applied to this particular reasoning problem, and from no other approach we are aware of. Our findings, presented in Table 1, confirmed these predictions.

Table 1: Percentage of conclusion types for Problems 1 and 2.

	Problem 1	Problem 2	Total
<i>A is taller than B B is taller than A</i>			
<i>A is taller than C C is taller than A</i>			
Single-subject conclusions	26%	45%	35%
Double-subject conclusions	14%	15%	14%
Nothing follows	54%	31%	43%
Other	6%	9%	8%

Another way to increase the relevance of a conclusion *A is more... than B and C* inferred from indeterminate relational premises is to act on the effect side. As we pointed out, a conclusion cannot yield more cognitive effect than the premises from which it is deductively derived. However, the information contained in the premises of a problem can yield greater or lesser cognitive effects, depending on the wider context. The greater are these effects, the more useful it is to derive a conclusion which is a step towards the production of these effects, and therefore the more relevant is this conclusion. Acting on the effect side here means providing or suggesting a context in which a conclusion derived from the premises of a problem might yield greater or lesser cognitive effects.

In Problem 2, the conclusion *A is taller than B and C* has a modicum of potential relevance. The cognitive effects that this conclusion might yield remain in the vague since no context is given. The relevance of such a conclusion can be increased by

manipulating the effect factor in the way we have just suggested. This can be done in particular by providing a context in which this conclusion will have clear contextual implications. Imagine for instance that the premises of Problem 2 are processed with the knowledge that the tallest person of A, B and C, is the tallest person in world. In this context, deducing that *A is taller than B and C* is a necessary step towards inferring that *A is the tallest person in the world*.

We predict that people should be more inclined to produce the conclusion *A is more... than B and C* when an appropriate context is given than when no context is given, or than when a less or non appropriate context is given. We tested this prediction in three experiments of an unpublished study done with Guy Politzer.

In the first experiment, participants received either a problem without explicit context (Problem 3) or a problem with an explicit context (Problem 4) and had to produce a conclusion:

Problem 3:

Premises: A is ahead of B
A is ahead of C

Problem 4:

Context: *A, B and C were the top three finishers in the race last Sunday.*

Premises: A is ahead of B
A is ahead of C

For both problems, it follows from the premises that *A is ahead of B and C*. However, in the race context, inferring the logical conclusion *A is ahead of B and C* is a step towards

inferring the contextual implication *A won the race*. The possibility of deriving this contextual implication endows the logical conclusion *A is ahead of B and C* with greater relevance than in the absence of any explicit context. Since the inference that A is ahead of B and C has a greater relevance in the race context, it should be more frequently performed, and participants should formulate more determinate conclusions and fewer “nothing follows” answers. Our results indeed show that Problem 4 resulted in a higher rate of determinate conclusions than Problem 3 (54% vs. 70%, $\chi^2(1) = 5.59$, $p < .02$). Moreover, in the race context, there were three times as many determinate conclusions referring to the race context like “*A is the first*” or “*A is the winner*” than conclusions simply integrating the two premises like “*A is ahead of B and C*” or “*B and C are behind A.*”

In a second experiment, we manipulated the effect factor by using two different explicit contexts both of which increased the relevance of the conclusion *A is more... than B and C*. However the context of Problem 5 (almost identical to that of Problem 4 above) produced a greater increase in relevance than that Problem 6:

Problem 5:

Context: *A, B and C were the first three finishers in the race last Sunday..*

Premises: A arrived before B
A arrived before C

Problem 6:

Context: *A, B and C were the last three three finishers in the race last Sunday.*

Premises: A arrived before B
A arrived before C

In Problem 5, the context explicitly focuses on people who were the first three in a race; if this is relevant at all, knowing who was *the* first should be even more relevant. The premises of the problem can thus achieve relevance by making it possible to infer who precisely arrived first and who did not. Deriving that A arrived before B and C enables one to infer three contextual implications: *A won the race, B did not win the race, and C did not win the race.* In Problem 6, the context focuses on people who arrived last in an athletics race. In contrast with Problem 5, deriving that A arrived before B and C makes it possible to infer only one contextual implication: *A did not arrive last.* Because the relation between B and C is indeterminate, it is impossible to infer who arrived last. The conclusion that A arrived before B and C has some relevance in Problem 6, but less so than in Problem 5 and should therefore be produced less often. Our results (see Table 2) show indeed that people derived more determinate conclusions in Problem 5 than in Problem 6 (94.4% vs. 74.7%, $\chi^2(1) = 13.45$, $p < .001$).

Table 2: Percentage of conclusion types for Problems 6 and 7

	Problem 5	Problem 6
	N = 90	N = 91
Determinate conclusions	94.4	74.7
Nothing follows	3.3	18.7
Errors and weird answers	2.2	6.6

Any explicit context evokes a wider *implicit* context of general knowledge. For instance, the explicit context of problem 5, “*A, B and C were the first three finishers in the race last Sunday,*” evokes background knowledge about racing, about the value attributed to winning, prizes or medals given to winners, and so on. So, inferring from the explicit context that A has won the race makes it possible to infer from the implicit context that A is likely to be pleased, that he may be given a medal or a prize, and so on.

In a third experiment, we manipulated relevance by evoking different implicit contexts. In general, when a context is explicitly provided, participants may expect the premises of a problem to be relevant in this explicit context or, at least, in the wider context implicitly evoked by this explicit context. If the explicit and implicit contexts are related in content to the premises, this should strengthen the expectation of relevance and encourage participants to derive positive conclusions from the premises rather than answer that nothing follows. Inversely, if the explicit and implicit contexts are unrelated in content to the premises, this should lower participant’s expectations of relevance and encourage them to say that nothing follows. Here is how we tested this prediction.

Consider Problems 7 and 8:

Problem 7:

Context: *A, B, and C, who were measured during a medical examination, are not of the same height*

Premises: A is taller than B
A is taller than C

Problem 8:

Context: *A, B, and C did not win the same amount of money at the last lottery,*

Premises: A is taller than B
A is taller than C

The explicit context of Problem 7, by mentioning measurements of height as part of a medical examination evokes an implicit context of common knowledge where differences in height may have implications for health, performance, accessibility to certain jobs, and so on. This should encourage participants to see the conclusion “*A is taller than B and C*” as potentially relevant in this implicit context. The explicit context of Problem 8, mentioning the winning of money in a lottery, evokes an implicit context of common knowledge where individual height plays no role at all. Hence, we should observe a much lower rate of determinate conclusions for Problem 8 than for Problem 7. Our results (see Table 3) confirmed that there were many more determinate conclusions for Problem 7 (76.5%) than for Problem 8 (42.9%, $\chi^2(1) = 38.8$, $p < .0001$).

Table 3: Percentage of conclusion types for Problems 8 and 9

	Problem 7	Problem 8
	N=162	N=168
Determinate conclusions	76.5	42.9
Nothing follows	19.8	50.6
Errors and weird answers	3.7	6.5

The experiments presented in this section give support to the cognitive principle of relevance, i.e. the claim that human cognition tends to be geared to the maximisation of relevance, by corroborating some of its consequences in the area of psychology of reasoning. More specifically, the choice to draw or not to draw conclusions from a given set of premises, and the choice of which particular conclusion to draw, if any, are guided by considerations of relevance. People are inclined to draw a specific conclusion from a set of premises to the extent that this conclusion seems potentially relevant. This is a non-trivial consequence of the cognitive principle of relevance. It has in turn non-trivial consequences for the study of reasoning in general. In particular, people's failure to derive some specific conclusion in a reasoning task may be due not to poor logical capacities or to pragmatic problems of comprehension of the task, but to the failure to see as relevant either the conclusion they were intended to draw, or, more subtly, the failure to see the relevance of some intermediary inferential step necessary for deriving the intended conclusion. In spontaneous inference, being guided by consideration of relevance should contribute to the overall efficiency of inferential processes, but it may

also, on occasion, prevent one from reaching some highly relevant conclusion because crucial intermediary steps didn't seem relevant at all.

4. Testing the Communicative Principle of Relevance with the Wason Selection Task

Task

Wason's Selection Task (Wason, 1966) has been the most commonly used tool in the psychology of reasoning (see Manktelow, 1999). Genuine versions of Wason's selection task share the same basic four-component structure:

- (i) An introduction (sometimes in a narrative form);
- (ii) A conditional statement known as the "rule", with the linguistic form "If P, then Q," and either a descriptive content stating how things are, or a deontic content stating how they should be;
- (iii) Four cards: one representing a case where P is satisfied, one where P is not satisfied, one where Q is satisfied, and one where Q is not satisfied (known respectively as the P, the not-P, the Q, and the not-Q cards). When the card displays information about P, information about Q is hidden, and conversely.
- (iv) The instruction to select all and only those cards where the hidden information must be made visible in order to judge whether the rule is true (in descriptive versions) or is being obeyed (in deontic versions).

For example the text of an "abstract" descriptive selection task might be: "Here are four cards. Each has a number on one side and a letter on the other side. Two of these cards are here with the letter side up, and two with the number side up. Indicate which of these cards you need to turn over in order to judge whether or not the following rule is true: "If there is a 6 on one side, there is an E on the other side".



With such an abstract version of the task, typically only about 10% of participants make the correct selection of the 6 and G cards, that is, the cards that represent the P case and the not-Q case.

In a typical example of a deontic version of the task (Griggs & Cox, 1982), participants are presented with a rule such as "If a person is drinking beer, then that person must be over 18 years of age", and with cards representing four individuals in a bar, with what they are drinking indicated on one side of the cards, and their age indicated on the other side. The four cards represent respectively a person drinking beer, a person drinking soda (with the age hidden for these first two persons), a person aged 29, and a person aged 16 (with the drink hidden for these two other persons). Participants are instructed to select the cards that must be turned over to see whether any of these four people is breaking the rule. Typically, the correct card combination (i.e., the P-card "This person is drinking beer" and the not-Q card "This person is 16 years old") is selected by well over 50% of the participants.

Work on the selection task has been the basis of a variety of claims about human reasoning and rationality. In particular, it has been taken to show that most individuals do not, in general, reason in accordance with the rules of logic, not even the elementary rules of propositional calculus, as evidenced by their failure to select the P and the not-Q cards in descriptive versions of the task (e.g., Cheng & Holyoak, 1985; Griggs & Cox, 1982). Does the selection task really provide a tool to test general claims about human

reasoning? Evans (1989) maintained that participants understand the task as one of identifying the *relevant* cards, and use, for this, heuristic cues of relevance rather than deductive reasoning. Extending this insight, Sperber, Cara and Girotto (1995) put forward a general explanation of the selection task based on relevance theory. They argued that participants' performance on the selection task is best explained by considering that (1) the very process of linguistic comprehension provides participants with intuitions of relevance, (2) these intuitions, just as comprehension generally, are highly content and context dependent, and (3) participants trust their intuitions of relevance and select cards accordingly. In standard versions of the task, these intuitions are misleading. In other versions, many deontic versions in particular, people's intuitions of relevance point towards the correct selection of cards. If, in the selection task, pragmatic comprehension mechanisms determine participants' response and thus pre-empt the use of whatever domain-general or domain-specific reasoning mechanisms people are endowed with, the task cannot be a good tool for the study of these reasoning mechanisms. On the other hand, it may be of some use in studying people's intuitions of relevance.

Participants presented with a Wason selection task approach the text of the problem, and in particular the conditional rule, in the same way in which they approach all utterances in conversation or in reading. They make use of their standard comprehension abilities. The very fact that a text is presented to them raises expectations of relevance, and they search for an interpretation that satisfies these expectations (which, given the artificiality of the task, may be quite modest). In doing so, they follow the relevance-guided comprehension procedure explained above in (8), that is, they follow a path of least effort in constructing interpretive hypotheses and stop when their expectations of

relevance are satisfied. This is, in particular, what participants do with the conditional rule of the selection task: guided by expectations of relevance, they derive from it consequences that might justify these expectations.

The rule itself, being a conditional statement, is not directly testable. Merely by looking at the two sides of a card, you can check the truth or falsity of a plain atomic statement or of a conjunction of atomic statements such as “there is a 6 on one side of this card and an E on the other side.” It is true if it matches your observations, and false otherwise. You cannot however confirm a conditional statement such as “if there is 6 on one side, then there is an E on the other side” by matching it to your observations. The truth of a conditional statement is tested indirectly, by deriving from it consequences that are directly testable and testing these. Participants have therefore two reasons to derive consequences from the rule. The first reason is to interpret it in a way that satisfies their expectation of relevance. The second reason is to find directly testable consequences of the rule in order to give a sensible response to the experimenter. What they do in practice is give a response that is based on the consequences they spontaneously derived in interpreting the rule, without looking for other consequences that might provide a better test of the rule. What they should do, in principle, is make sure that not only the consequences they derive are entailed by the rule but also that, conversely, the rule is entailed by these consequences. Otherwise, the consequences might be true and the rule false. This would involve more than just reasoning in accordance with the rules of propositional calculus. It would require also higher-order reasoning about the structure of the problem. People’s failure to do so shows not that, presented with such a problem, they are illogical, but that they are unreflective or, at least, insufficiently reflective, and overconfident in their intuitions of relevance.

In the case of the abstract task described above, participants may infer from the rule “If there is a 6 on one side, there is an E on the other side” that the card with 6 must have an E on the other side. They may also infer from the rule the consequence that there are cards with a 6 and an E (otherwise the rule would be irrelevant). Making either or both of these consequences part of the interpretation of the rule contributes to its relevance by indicating what one might expect to see when turning over the cards. If participants use the first of these two consequences to decide which cards must be turned over in order to see whether the rule is true or false, they will select just the card with a 6 (the P card). If they use just the second consequence, or if they use both, they will turn over the card with a 6 and that with an E (the P card and the Q card). These are indeed the most frequent selections with standard selection tasks. In a deontic case such as that of the drinking-age problem, participants might, in order to satisfy their expectations of relevance, derive from the rule (“If a person is drinking beer, then that person must be over 18 years of age”) the consequence that there should be no beer drinker under 18. They would then select the card representing a beer drinker (the P card) and that representing a person under 18 (the not-Q card), thus, as it happens, providing the correct selection.

Why should the consequences derived in the two problems be different? Because they are derived in their order of accessibility until expectations of relevance are reached, and both order of accessibility and expected level of relevance are context-dependent. In both problems—and in general with conditional statements—, the most accessible consequence is the *modus ponens* one: in the abstract problem, it is that the card with a 6 should have an E on the other side, and in the drinking age problem, it is that the beer drinker should be 18 or above. In both cases, this implication determines the selection of the P card, which is indeed selected by most participants in both experiments. Why, then,

do many participants select also the Q card in the abstract version, and the majority of participants select the not-Q card in the drinking age problem (as in most deontic versions of the task)? In the abstract problem above, the implication “there are cards with a vowel and an even number” is much more easily accessed than the implication “there are no cards with a vowel and without an even number,” and satisfies the low expectations of relevance raised by this artificial problem. In the drinking age problem, on the other hand, the implication that there should not be underage beer drinkers is the most accessible and the only one that satisfies expectation of relevance: commonsensically, the point of a normative rule such as “If a person is drinking beer, then that person must be over 18” is not to make adult beer drinkers more common, but to make underage beer drinker less common.

By devising appropriate rule-context pairs, the order of accessibility of consequences and expectations of relevance can be manipulated and it should be possible to elicit different pattern of selection, including logically correct selections. Sperber et al. (1995) produced several *descriptive* versions of the task that elicited a higher percentage of correct responses than had ever been found before with such versions. They showed that – contrary to what was generally believed at the time – good performance is not restricted to deontic versions.⁶ Girotto, Kimmelmair, Sperber & Van der Henst (2001) provided further evidence for the relevance approach by demonstrating how it can be used to manipulate *deontic* versions of the task and obtain at will either the common correct P and not-Q selections or incorrect P and Q selections (more commonly found in descriptive versions). Further experiments and comparisons with the approach of Leda Cosmides and

⁶ Other studies have confirmed this: e.g., Green & Larking, 1995; Hardman, 1998; Johnson-Laird & Byrne, 1995; Liberman & Klar, 1996; Love & Kessler, 1995.

her collaborators (Cosmides, 1989; Fiddick, Cosmides & Tooby, 2000) can be found in Sperber & Girotto (2002, in press). Here, by way of illustration, we give just two examples of these experiments, one succinctly, the other in greater detail.

Girotto, et al (2001) used the following problem (adapted a problem from Cheng & Holyoak, 1985): “Imagine that you work in a travel agency and that the boss asks you to check that the clients of the agency had obeyed the rule ‘If a person travels to any East African country, then that person must be immunized against cholera,’ by examining cards representing these clients, their destinations and their immunizations.” The four cards indicated “Mr. Neri. Destination: Ethiopia”, “Mr. Verdi. Destination: Canada”, “Immunizations done: Cholera” and “Immunizations done: None”, respectively, and as usual, participants were asked which card had to be examined in order to find out whether the rule had been obeyed by the clients of the agency. In this context, the relevance of the rule is to prevent people without cholera immunization from travelling to East African countries. We predicted therefore that participants would choose the P card (a traveller to an East African country) and the not-Q card (a person without cholera immunization). Such a prediction is not specific to relevance theory. It would be shared by all researchers in the area, whatever their theoretical viewpoint. It reiterates, after all, common findings, that have been explained, for instance, by proposing that people have pragmatic reasoning schemas for reasoning about obligations and permission (Cheng & Holyoak, 1985), or that they have an evolved “Darwinian algorithm” for reasoning about social contracts (Cosmides, 1989).

According to the relevance-theoretic approach, what causes the selection of the P and not-Q cards in this deontic scenario is that the presence of individuals violating the cholera rule among the people represented by the cards would be more relevant than the

presence of individuals obeying the rule. Could this relative relevance of cases of violation vs. cases of conformity be reversed by altering the context, which, if the relevance approach is correct, should cause participants to choose the P and the Q card? To do this, we used the same scenario, with a twist. The narrative stated that contrary to what the boss of the agency had thought, cholera immunization is not required anymore when travelling to East Africa. The boss is now worried that she may have misinformed clients and caused them to follow a rule that is no longer in force. She then asks the employee to see whether or not clients have obeyed the rule “If a person travels to any East African country, then that person must be immunized against cholera” by looking at cards similar to those used in the previous condition. In this context, what is relevant is that some clients may have followed the false rule and that they may have been immunized unnecessarily (and might, for instance, sue the agency). On the other hand, the case of clients who have ignored the rule is no longer relevant. We predicted therefore that participants would select the P card (a traveller to an East African country) and the Q card (a person with cholera immunization). Note that this prediction is non-standard but follows from the relevance-based explanation of the selection task. This prediction was confirmed. Table 4 shows the results we obtained in a within participants design. (We also obtained practically the same results with a between-subject design.)

Table 4: Percentage of the main selection patterns in the true and false cholera rule selection task.

Pattern	True rule	False rule
P and not-Q	62	15
P and Q	26	71
Other	12	14

This cholera-rule experiment gives, we hope, an intuitively clear illustration of the role of relevance in participants' response to selection task problems. However, it remains too intuitive to give a truly specific confirmation to the communicative principle of relevance. In particular, it throws no light on the respective role of effect and effort in guiding participants' intuitions of relevance and selection of cards.

In their Experiment 4, Sperber et al. (1995) aimed at taking apart the two factors of relevance, effect and effort, testing their respective roles, and ascertaining whether relevance, which combines the effort and the effect factors in a principled manner, is more explanatory than effort or effect taken alone. For this, they created four scenarios, varying the effect and the effort factors separately in four conditions names **effect-/effort+**, **effect-/effort-**, **effect+/effort+**, and **effect+/effort-** (see Figure 1). All four scenarios involved a machine that manufactures cards with a number on one side and a letter on the other side. A character, Mr. Bianchi, asserts: "If a card has a 6 on the front, it has an E on the back." In all conditions, the four cards had respectively a 6, a 4, an E and an A on the visible side, and participants were asked which card or cards had to be turned over to check whether what Mr. Bianchi says is true.

Figure 1: The four conditions of the machine experiment (Sperber et al. 1995).

<i>Effect-/Effort +</i>	<i>Effect-/Effort-</i>	<i>Effect+/Effort+</i>	<i>Effect+/Effort -</i>
<p>A machine manufactures cards.</p> <p>It is programmed to print at random, on the front of each card,</p>			
A number	a 4 or a 6	A number	A 4 or a 6
On the back of each card, it prints a letter at random.	On the back of each card, it prints either an E or an A at random.	On the back of each card, it prints a letter: - When there is a 6, it prints an E. - When there is not a 6, it prints a letter at random.	On the back of each card, it prints a letter: - When there is a 6, it prints an E. - When there is a 4, it prints an E or an A at random.
		One day, Mr Bianchi, the person in charge, realises that the machine has produced some cards it should not have printed. On the back of the cards with a 6, the machine has not always printed an E: sometimes it has printed any letter at random.	sometimes it has printed an A instead of an E.
The person in charge, Mr Bianchi, examines the cards and has the strong impression that the machine does not really print letters and numbers at random. I think, he says, that		Mr Bianchi fixes the machine, examines the newly printed cards and says: don't worry, the machine works fine,	
if a card has a 6 on the front, it has an E on the back			

From the conditional “If a card has a 6 on the front, it has an E on the back” participants are sure to derive consequence (12). They may also derive either or both of (13) and (14):

(12) The card with a 6 has an E on the other side

(13) There are cards with a 6 and an E

(14) There are no cards with a 6 and without an E

In the two **effort+** conditions, (13) is easier to derive than (14), which involves two negations. Moreover (14) does not carry any obvious effect worth the extra effort. So we should expect participants to base their selections either on (12) and to select just the E, or on (12) and (13), and to select both the E and the 6.

To increase the probability that participants would derive consequence (14) before (13), we could act on the effort side or on the effect side. To act on the effort side, we had, in the two **effort-** conditions, the machine print only 6s and 4s on one side and Es and As on the other side. Instead of an indefinite number of possible number-letters combinations (e.g. $9 \times 26 = 234$ if only numbers from 1 to 9 are used), we have now four possible combinations; 6 and E, 6 and A, 4 and E, and 4 and A, which are all equally easy to represent. This makes it possible to simplify (14) and replace it with (14')

(14') There are no cards with a 6 and an A


We predicted that (14') being easier to represent than (14), more participants would derive it and would, accordingly, select the card with an A rather than the card with an E in the **effort-** conditions than in the **effort+** conditions.

To increase the probability that participants' expectations of effect would be satisfied with an interpretation of the rule as implying (14) rather than (13), we developed the scenario, in the two **effect+** conditions, as follows: the machine was supposed to print an E on the back of cards with a 6; however the machine ceased to function properly and printed cards with a 6 and a letter other than an E; after having repaired it, Mr. Bianchi, asserted: "If a card has a 6 on the front, it has an E on the back." In such a context, the relevance of Mr. Bianchi's assertion went through the implication that there were no cards with a 6 and a letter other than an E (in other terms, consequence (14)). On the other hand, in such a context, consequence (13) does not contribute to the relevance of the conditional. We predicted therefore that in the two **effect+** conditions, participants would more often infer (14) and select the 6 and the A card than in the **effect-** conditions.

The two **effect+** conditions on the one hand, and the two **effect-** conditions on the other hand, differ from one another only on the effort side, while the two **effort+** and the two **effort-** conditions differ from one another only on the effect side. Given this, the predictions that follow from the relevance-theoretic account of the task are self-evident: the best performance should be with the **effect+/effort-** condition, and the worse one with the **effect-/effort+** condition. The performance on the **effect+/effort+** and on the **effect-/effort-** condition should be at an intermediary level between the two other conditions. Moreover, the two factors, effect and effort, should, each on its own, contribute to good performance.

The results are summarised in figure 2:

Figure 2: Percentage of 6 and E (incorrect) and 6 and A (correct) responses in the four versions of the Machine problem.



These results confirm our prediction. Both factors of relevance, effect and effort, were shown to play a role in performance. These results show how effort and effect factors can be manipulated independently or jointly so as to favour one interpretation of a conditional statement over another. The advantage of the selection task paradigm in this context is that participants' interpretations of the rule are rendered manifest by their selection of cards.

5. Testing the Communicative Principle of Relevance with a speech production task

According to the communicative principle of relevance, utterances convey a presumption of their own optimal relevance, and do so whether or not they actually are optimally relevant. Speakers may fail to achieve relevance, or they may not even try, and, in such cases, the presumption of optimal relevance is unjustified. Justified or not, it is automatically conveyed by every utterance used in communication, and it guides the process of comprehension. Most research exploring the consequences of the communicative principle of relevance have, accordingly, focused on the comprehension process. Still, the communicative principle could not be right—and relevance could not guide comprehension—if speakers were not, often enough, trying to be optimally relevant, and successful at it. In the study that we report in this section, we investigate the degree to which speakers actually aim at being relevant, even when talking to perfect strangers from whom they have little to expect in return.

Imagine the following exchange between two strangers in the street:

(15) Mr X: Hello, do you have the time, please?

Mrs Y: Oh yes, it is 4:30

In fact, Mrs Y's watch does not indicate 4:30 but 4:28. She has chosen to round her answer even though she could have been more accurate. Rounding numbers is quite common. People round when talking about money, distance, time, weight, and so on. What explains this behaviour? We recently proposed that rounding is in part explained by considerations of relevance (Van der Henst, Carles & Sperber, 2002). A rounded

answer is generally more relevant than an accurate one, and speakers round in order to be relevant to their hearer.

In a few situations, when taking a train for instance, a person asking for the time is better off with an answer precise to the minute. If your train leaves at 4:29, and you are told that it is 4:30 while it is in fact 4:28, you may believe that you've missed it when in fact you could still catch it. On the other hand, if you were told that it is 4:25, you might end up missing your train by considering that you still had four minutes to board it. In most situations, however, the consequences you would draw from a time rounded to the nearest multiple of five minutes are the same as those you would draw from a time accurate to the minute. So, in general, rounding does no harm. Does it do any good? Rounded number requires less *processing effort*. 4:30 is easier to manipulate than 4:28. Communicating rounded numbers may thus be a way to provide an optimally relevant answer to addressees by reducing their processing effort without compromising any cognitive effect likely to be derived.

In most situations then, a speaker who is asked for the time and wishing to be as relevant as possible would round her answer. She might however be rounding for other reasons. In particular, if she wears an analogue watch indicating only numbers that are multiple of five, it may be easier for her to round than not to round. She might then round to minimise, not her audience's effort, but her own. In fact, a sceptic might argue, the goal of minimising one's audience's effort might not play any role in the tendency of people asked for the time to give a rounder answer.


In order to find out whether a tendency to optimize relevance was a factor in rounding the time, we approached people on the campus of the Paris University of Paris VII and just asked them: "*Hello, Do you have the time please?*" (Van der Henst, Carles

& Sperber, 2002). We took note of their response and of the type of watch they were wearing: analogue or digital, and distinguished two groups, the “analogue” and the “digital” group. For people with a digital watch, it requires less effort to just read aloud the exact time indicated by their watch than to round it to the closest to multiple of five. If people asked for the time were just trying to minimize their own effort, then they should always round when their watch is analogue, and never do so when it is digital. On the other hand, if people are also motivated by the goal of reducing their audience’s effort, then, not only people with analogue watches, but also a significant percentage of people with digital watch should round.

What we found is that people rounded in both conditions. The percentage of rounders is calculated on the basis of the percentages of responses which indicate the time in a multiple of five minutes. If people never rounded there should be 20% of such responses (this is the theoretical distribution of numbers which are multiples of 5). However, the percentages we observed in the two conditions were much higher: 98% of answers were multiple of 5 in the analogue group, and 65.8% in the digital group. This means that 97% of people rounded in the analogue condition and 57% in the digital one⁷ (see Figure 3). Hence, even though participants of the digital group rounded less than participants of the analogue group, a majority of them did, remarkably, made an extra effort in order to diminish the effort of their audience.

⁷ To calculate the percentage of rounders we used the following formula: *Percentage of rounders* = $(M - 20)/80$, where M is equal to the percentage of answers given in a multiple of five. When M is equal to 20, the percentage of rounders is equal to 0, when it is equal to 100, so is the percentage of rounders.

Figure 3: Percentages of rounders in the three experiments: In Experiment 1, participants wore analogue watches in the “analogue” group, and digital watches in the “digital” group; in Experiment 2, participants were just asked for the time in the “control” group, and were asked for the time by an experimenter setting his watch in the “experimental” group; in Experiment 3, participants were asked for the time more than 15 minutes before the time at which the experimenter said he or she had an appointment in the “earlier” group, and less than 15 minutes before the appointment in the “later” group.



Some people with analogue watches may round just in order to save their own effort, but the case of people with digital watch shows that a majority of people are disposed to round, even when this means making an extra effort. We attributed this disposition to a more general disposition, that of trying to produce optimally relevant utterances. Still, an alternative explanation could be that people round in order to minimize their commitment: they may not be sure that their watch is precise to the minute, and be more confident that it is accurate within a five-minute interval. Indeed, this desire to minimize commitment may account for some of the rounding we observed,

but could it be enough to make the relevance-based explanation superfluous? To investigate this possibility, we created a situation where accuracy manifestly contributed to relevance.

Although rounded answers are easier to process than non-rounded ones, there are some situations, such as that of the train evoked above, where optimal relevance depends upon cognitive effects that are carried only by a more accurate answer. Speakers guided by the goal of producing an optimally relevant answer should, in this condition, provide, if they can, a more precise answer than in the ordinary kind of situation in which our first experiment took place.

We tested this prediction in Experiment 2 with two groups of people. In the control group, participants were approached in the same way as in the previous experiment and were just asked for the time. In the experimental group, the request for the time was framed in a context in which an accurate answer was obviously more relevant. The experimenter approached the participant with a watch held in his hand and said: *«Hello! My watch isn't working properly. Do you have the time please?»* In this context, it was clear that the experimenter was asking for the time in order to set his own watch and that, for this purpose, an answer precise to the minute would be more relevant. Only the answers of participants with an analogue watch were recorded. Participants had therefore to make an extra effort in order to provide an accurate answer. We found that participants were much more accurate in the experimental than in the control condition: there were 94% of rounders in the control condition and only 49% in the experimental one (see Figure 3, Experiment 2). This means that 51% of participants of the experimental group gave the requester a time accurate to the minute. Note that rounded answers may nevertheless have been in conformity with the presumption of optimal

relevance: even if approximate, they were relevant enough to be worth the hearer's attention, as required by the first clause of the presumption, and, as required by second clause, they may have been the most relevant ones compatible with the speakers' abilities (if they had doubts about the accuracy of their watch), or preferences (if they were reluctant to work out a more precise answer). Our results show anyhow that a majority of the people not only understood that accuracy was more relevant in this condition but also were able and willing to make the effort of giving an accurate answer.

That accuracy to the minute is relevant to someone setting his watch is easy enough to understand. It need not involve the kind of refined concern for relevance that relevance theory presupposes. In a third experiment, we manipulated the relationship between relevance and accuracy a much subtler way.

Suppose you want to know how much time you have left before an appointment at 4:00pm. The more you will approach the time of the appointment the more accuracy is likely to be relevant. At 3:32, being told that it is 3:30 is likely to have practically the same effect as being told, more accurately, that it is 3:32. On the other hand being told at 3:58 that it is 4:00, is likely to be misleading. Two minutes may, for instance, be the time you need to reach the place of your appointment. In other words, the closer you are to the time of the appointment, the more accuracy becomes relevant.

In the third experiment, all participants were approached in the same way and told *"Hello Do you have the time please? I have an appointment at T"* We then divided participants into two groups: the "earlier" group of those who gave a time between 30 to 16 minutes before the time of the appointment and a "later" group of those who answered with a time between 14 minutes before the time of the appointment to the time of the appointment itself. As we had predicted, the results show that participants rounded

less in the “later” group (75% of participants) than in the “earlier” group (97%). 22% difference may not seem so impressive until you realize that those people in the later group who did give an accurate answer, not only were willing to make the effort of reading more carefully their analogue watch and had enough confidence in its accuracy, but also made the extra effort of taking the perspective of the stranger who was addressing them and of inferring that accuracy, at this point in time, would contribute to the relevance of their utterance.

The experiments described in this section show how subtle aspects of people’s spontaneous speech behaviour can be predicted on the basis of the communicative principle of relevance: speakers tend to produce utterances that justify the presumption of optimal relevance these utterances automatically convey.

6. Conclusion

The studies reported in this chapter tested and confirmed predictions directly inspired by central tenets of relevance theory and, in particular, by the cognitive and the communicative principles of relevance. Of course, it would take many more successful experiments involving a variety of aspects of cognition and communication to come anywhere near a compelling experimental corroboration of relevance theory itself. Still, from a pragmatic point of view, the few experiments we have presented here, together with others we have mentioned, show, we hope, how imagining, designing and carrying out experiments helps expand and sharpen pragmatic theory. From an experimental psychology point of view, these experiments illustrate how a pragmatic theory that is precise enough to have testable consequences can put previous experimental research in a novel perspective and can suggest new experimental paradigms.

References

- Braine, M.D.S., & O'Brien, D.P. (1998). *Mental Logic*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Byrne, R.M.J. & Johnson-Laird, P.N. (1989). Spatial reasoning. *Journal of Memory and Language*, 28, 564-575.
- Cheng, P.N., & Holyoak, K.J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 391-416.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, 31, 187-276.
- Evans, J.St.B.T. (1989). *Bias in human reasoning : Causes and consequences*. Hove, UK : Lawrence Erlbaum Associates.
- Evans, J.St.B.T., Newstead, S.E., & Byrne, R.M.J. (1993). *Human reasoning. The Psychology of Deduction*. Hove, UK: Lawrence Erlbaum Associates.
- Fiddick, L., Cosmides, L., & Tooby, J. (2000). No interpretation without representation: The role of domain-specific representations in the Wason selection task. *Cognition*, 77, 1-79.
- Gibbs, R.W. and J.F. Moise (1997). Pragmatics in understanding what is said. *Cognition* 62, 51-74
- Giroto, V., Kimmelmair, M., Sperber, D., & van der Henst, J.B. (2001). Inept reasoners or pragmatic virtuosos? Relevance and the deontic selection task. *Cognition*, 81, 69-76.
- Green, D. W., & Larking, R. (1995). The locus of facilitation in the abstract selection task. *Thinking and Reasoning*, 1, 183-199.

- Griggs, R.A., & Cox, J.R. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, 73, 407-420.
- Happé, F. (1993) "Communicative competence and theory of mind in autism: a test of relevance theory." *Cognition*, 48, 101-119.
- Hardman, D. (1998). Does reasoning occur in the selection task? A comparison of relevance-based theories. *Thinking and Reasoning*, 4, 353-376.
- Harman, G. (1995). Rationality. In Smith, E.E. Oherson, D.N, (Eds.) *Thinking: An invitation to cognitive science*, Vol. 3 (2nd Ed.)
- Johnson-Laird, P.N. & Byrne, R.J.M. (1991). *Deduction*. Hove, U.K: Lawrence Erlbaum Associates.
- Johnson-Laird, P.N. & Byrne, R.J.M. (1991). *Deduction*. Hove, U.K: Lawrence Erlbaum Associates.
- Johnson-Laird, P.N., & Byrne, R.M.J. (1995). A model point of view. *Thinking & Reasoning*, 1, 339-350.
- Jorgensen, J., G. Miller and D. Sperber (1984) "Test of the mention theory of irony." *Journal of Experimental Psychology: General*, 113, 112-120
- Lieberman, N., & Klar, Y. (1996). Hypothesis testing in Wason's selection task: Social exchange, cheating detection or task understanding. *Cognition*, 58, 127-56
- Love, R., & Kessler, C. (1995). Focussing in Wason's selection task: Content and instruction effects. *Thinking and Reasoning*, 1, 153-82
- Manktelow, K.I. (1999). *Reasoning and Thinking*. Hove: Psychology Press.
- Matsui, T. (2000). *Bridging and Relevance*. John Benjamins, Amsterdam

- Matsui, T. (2001). Experimental Pragmatics: Towards Testing Relevance-Based Predictions about Anaphoric Bridging Inferences. In V. Akman et al. (Eds.), *Context 2001*, Berlin: Springer-Verlag, 248–260.
- Medin, D.L., Coley, J.D., Storms, G. & Hayes, B. (In Press). A Relevance Theory of induction. *Psychonomic Bulletin and Review*.
- Nicolle, S. & Clark, B. 1999: Experimental pragmatics and what is said: A response to Gibbs and Moise. *Cognition* 66, 337-354.
- Noveck, I. A. (2001). When children are more logical than adults: Investigations of scalar implicature. *Cognition*, 78, 165-188.
- Noveck, I. A. & Posada, A. (2003). Characterizing the time course of an implicature: An evoked potentials study. *Brain and Language*, 85, 203-210.
- Noveck, I. A., Bianco, M., & Castry, A. (2001). The costs and benefits of metaphor. *Metaphor and Symbol*, 16, 109-121.
- Poltzer, G. (1996). A pragmatic account of a presuppositional effect. *Journal of Psycholinguistic Research*, 25, 543-551.
- Rips, L.J. (1994). *The psychology of proof*. London: M.I.T Press.
- Ryder, N & Leinonen, E. (2003). Use of Context in Question Answering by 3-, 4- and 5-Year-Old Children. *Journal of Psycholinguistic Research*, 32, 397-415

- Sperber, D. (1994). 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*, New York : Cambridge University Press, 39-67.
- Sperber, D., & Girotto, V. (2002). Use or misuse of the selection task? Rejoinder to Fiddick, Cosmides, and Tooby. *Cognition* 85, 277-290.
- Sperber, D., & Girotto, V. (in press). Does the Selection Task Detect Cheater-Detection? In Fitness, J. & Sterelny, K. (eds.), *New directions in evolutionary psychology*, Macquarie Monographs in Cognitive Science, Psychology Press.
- Sperber, D., & Wilson, D. (1981). Irony and the use-mention distinction. In Peter Cole (ed.) *Radical Pragmatics*, 295-318. New York, NY: Academic Press
- Sperber, D., & Wilson, D (2002). Pragmatics, Modularity and Mind-reading. *Mind and language*, 17, 3-23)
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and Cognition. Second edition*. Oxford: Blackwell
- Sperber, D., Cara, F., & Girotto, V. (1995). Relevance theory explains the selection task. *Cognition*, 52, 3-39.
- Van der Henst, J.B., Carles, L & Sperber, D. (2002). Truthfulness and relevance in telling the time. *Mind and Language*, 17, 457-466.
- Van der Henst, J.B., Sperber, D. & Politzer, G. (2002) When is a conclusion worth deriving? A relevance-based analysis of indeterminate relational problems. *Thinking and Reasoning*, 8, 1-20.

Wason, P.C. (1966). Reasoning. In B. M. Foss. (Ed.) *New Horizons in Psychology*.
Harmondsworth: Penguin.

