Mind & Society (2006) 5:229–245
DOI 10.1007/s11299-006-0019-x

ORIGINAL ARTICLE

Symposium on "Cognition and Rationality: Part I" Relevance effects in reasoning

Jean-Baptiste Van der Henst

Received: 3 June 2004 / Accepted: 3 February 2005 /
Published online: 30 August 2006

Abstract Reasoning research has focussed mainly on the type of cognitive processes involved when representing premises and when producing conclusions. But less is known about the factors that guide these representational and inferential processes. What premises are actually taken as input in reasoning? And what conclusions are intended? In this paper it is argued that considerations of relevance (Sperber and Wilson, Relevance: communication and cognition. Blackwell, Oxford, 1995) are helpful for addressing these issues as a pragmatic analysis of two sorts of tasks is carried out, Wason's 2-4-6 problem (Study 1) and a conditional reasoning problem (Study 2). Study 1 indicates that the way this task is communicated may encourage participants to consider misleading information as highly relevant for solving it. Two experiments go on to show that when the relevance of misleading information is contextually diminished, participants are more efficient at providing the correct solution. Study 2 compares the production rate of two sorts of conclusions: logically valid but weakly relevant conclusions and invalid but relevant and pragmatically justified conclusions. This study shows that the relevance of conclusions determines to a large extent whether or not they will be produced.

Keywords Relevance and reasoning · 2-4-6 Problem · The relevance of conclusions

1 Introduction

Reasoning is often characterised by experimental psychologists as a cognitive activity that consists in inferring conclusions on the basis of premises. Con-

J.-B. Van der Henst (✉)
Institut des Sciences Cognitives, 67, Boulevard Pinel, 69675 Bron cedex, France
e-mail: vanderhenst@isc.cnrs.fr

sequently, two subgoals have to be achieved by our cognitive system when we reason. The premises must be represented in a format that can fit with a natural inferential procedure and this procedure must be applied to the represented premises in order to draw a conclusion. In line with these two subgoals, research on human reasoning addresses two questions: (1) How is the input information (i.e., the premises) represented? and (2) How is the output information (i.e., the conclusion) derived?

The nature of representational and inferential processes is a matter of debate among reasoning researchers. Some think that input information is represented by propositions and that mental logical rules are applied to these propositions in order to reach conclusions (Clark 1969; Rips 1994; Braine and O'Brien 1998). Others claim that input information is analogically represented by mental models on the basis of which conclusions can be read off and assessed (Huttenlocher 1968; Johnson-Laird 1983; Johnson-Laird and Byrne 1991). Consider for instance the premises *Pete is to the left of Dick and Dick to the left of Bill*. They can be represented by the propositions LEFT (Pete, Dick), LEFT (Dick, Bill). Applying a transitivity rule [i.e., LEFT (X, Y) & LEFT (Y, Z) –> LEFT (X, Z)] to these premises can result in the derivation of the conclusion *Pete is to the left of Bill.* Alternatively, the premises can give rise to the building of a mental model from which the relation between Pete and John emerges:

Pete Dick Bill

Answering the two *How* questions by analysing the mechanisms underlying the representation of premises and the derivation of conclusions leads one to address two *What* questions: (1) What information is taken as input for the representational processes? (2) What information emerges from the output of inferential processes? The type of reasoning processes certainly determines to some extent what is reasoned about since people cannot represent information and cannot draw conclusions that surpass the capacity of these processes. Take for instance the principle of truth which is one of the core principles of mental model theory (Johnson-Laird and Savary 1999; Johnson-Laird and Byrne 2002). It stipulates that due to working memory limitations, reasoners tend to build mental models that represent only what is true but not what is false. This implies that the false possibilities are not represented and that within the true possibilities only the true clauses tend to be represented. For instance, the exclusive disjunction ''there is king or else there is an ace'' is likely to be represented with the two models:

King
Ace

This set of models does not include the false possibility that there is neither a king nor an ace or the false possibility that there is both a king and an ace. Moreover the model on the first line does not explicitly represent the false

clause that there is not an ace and the model on the second line does not explicitly represent the false clause that there is not a king. A more explicit representation would involve the two following models:

King	¬Ace
¬King	Ace

As shown by Johnson-Laird and his colleagues the failure to represent what is false lead reasoners to very compelling but invalid inferences called *illusory inferences* (Johnson-Laird and Savary 1999). Another example concerns the inferential processes as they are described by the mental logic approach. According to this framework, human beings are equipped with a mental rule that corresponds to the *Modus Ponens* inference: *If P then Q*; *P*; therefore *Q*. In contrast, there is no such rule for the *Modus Tollens* inference: *If P then Q*; *not-Q*; therefore *not-P*. It follows that reasoners are more likely to draw the conclusion "*Q*" from "*If P then Q*" and "*P*" than the conclusion "*not-P*" from "*If P then Q*" and "not-*Q*". In brief, the very nature of our reasoning processes restricts the information we can represent and derive.

However, the restrictions imposed on the *What* questions by the *How* questions are unlikely to properly mark out the piece of information reasoners will actually process and derive. In experimental settings as well as in daily life we do not necessarily use the whole set of information that is made available (by the experimenter or the environment). We may consider that some premises and conclusions will be more interesting and more important than others and, given the limitation of our cognitive capacities, we may tend to devote time and effort to information from which we could benefit. In other words, there might be premises and conclusions that we will not represent and not derive, not because they are beyond the scope of our representational and inferential processes but because other premises and conclusions are preferable. Hence, in order to better understanding human reasoning, it might be helpful to analyse the factors that guide the selection and the derivation of information when we reason (apart from the capacity of inferential processes).

In my view, considerations of relevance are helpful for addressing these issues and the goal of this paper is to show the existence of relevance effects in the selection of input information and in the production of output information (see Fig. 1). The concept of relevance I will use is the one developed by Dan Sperber and Deirdre Wilson (Sperber and Wilson 1995; for another concept of

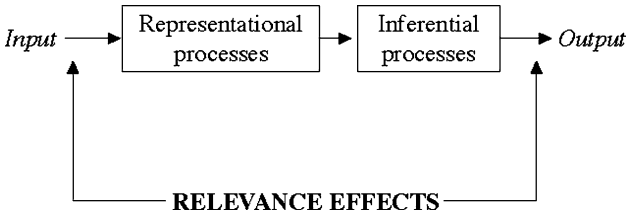

Fig. 1 Relevance effects in reasoning

relevance, see Evans 1995, and for an evaluation of the relevance approaches in the domain of reasoning see Hardman 1998). In the next section, I will briefly present Relevance theory. In a subsequent section, I will show how performance on the 2-4-6 task (Wason 1960) is partially determined by relevance effects in the selection of input information. Finally, I will describe a conditional reasoning study displaying relevance effects in the production of conclusions.

2 Relevance theory

Sperber and Wilson (1995) characterise the concept of relevance as a property of cognitive inputs. An input is relevant to individual when it carries *cognitive effects* in the context it is processed. The concept of cognitive effect captures the intuitive notion of improvement of knowledge; e.g., when a piece of information allows an individual to increase or refine her knowledge on a topic that is important to her. Contextual implications (i.e., new beliefs that can only be deduced from the input and the context taken together) are prime examples of cognitive effects as are the reinforcement of, the correction of or the rejection of previously held beliefs. The degree of relevance of an input relies on the amount of cognitive effects one may achieve by processing that input. This does not mean however that an absolute measure of relevance can be attained. Relevance is a comparative notion. All other things being equal, if processing stimulus A results in more cognitive effects than processing stimulus B, then A will be more relevant than B.

The production of cognitive effects is not the only factor that contributes to determining the degree of relevance of a piece of information. Another concerns the cognitive resources that are required. That is, two inputs conveying the same cognitive effect may differ in terms of *cognitive effort*: the cognitive effect an input carries may be achieved at a greater cost than the other. When effort increases, the sought improvement of knowledge becomes difficult to attain. The amount of effort required to process an input will therefore be an inverse function of its relevance. Hence, the degree of relevance relies on two factors—cognitive effects and cognitive effort.

2.1 Relevance of an input to an individual

(a) All other things being equal, the greater the positive effects achieved by processing an input, the greater the relevance of the input to the individual at that time.
(b) All other things being equal, the greater the processing effort expended, the lower the relevance of the input to the individual at that time.

Let's explain by example. Hugo wants to buy a new computer with the money he has left in his bank account. He arrives at the bank and inquires about the

amount of money in his account. Here are three alternative answers that the teller could give him:

[1] You have 2,200 euros in your account
[2] You have between 1 euro and 1,000,000 euros in your account
[3] The amount of money you have on your account is equivalent in Venezuelan Bolivars to the root cube of this number: 1297126029226368810 60,928

These answers differ with respect to their degree of relevance. [1] Is more relevant than [2] because it conveys more cognitive effects. From [2], Hugo can only infer that he is not overdrawn and that therefore the possibility of buying a computer is open. In contrast, from [1], Hugo knows that he can buy any computer that does not cost more than 2,200 euros. [1] is also more relevant than [3]: both statements convey the same information about the amount of money on Hugo's account and therefore yield the same cognitive effects but [3] requires much more processing effort than [1].

Sperber and Wilson argue that our cognitive mechanisms tend to operate in an efficient way, meaning by maximizing the benefits and minimizing the costs. The idea of cognitive efficiency is at the heart of the *cognitive principle of relevance which stipulates that ''human cognition tends to be geared towards the maximization of relevance*''. In other words, cognitive resources are likely to be allocated to the processing of the most relevant information available (see p. 261 in Sperber and Wilson 1995). According to Sperber and Wilson, the tendency of maximizing relevance is automatic and is the product of a selective pressure mechanism towards efficiency: the human cognitive system has evolved ''in such a way that our perceptual mechanisms tend automatically to pick out potentially relevant stimuli, our memory retrieval mechanisms tend automatically to activate potentially relevant assumptions, and our inferential mechanisms tend spontaneously to process them in the most productive way'' (Wilson and Sperber 2004).

Communication takes advantage of this tendency of maximizing relevance. A communicated piece of information is an ostensive stimulus. It aims to attract the audience's attention to a certain content and thus raises expectations of relevance. This leads to an important distinction between information obtained by communication and information not obtained by communication. The communicator informs the audience about his intention to inform the audience. He wants his audience to pay attention to the piece of information he communicates because it is relevant to the audience. The second, or communicative, principle of relevance stipulates that *''every act of ostensive communication communicates a presumption of its own optimal relevance*''. An ostensive act of communication is optimally relevant when (1) it is relevant enough to be worth being processed by the audience, (2) it is the most relevant one compatible with the communicator's abilities and preferences. In other words, an optimally relevant stimulus must reach a minimum level of relevance—the effects being sufficiently high and the effort being sufficiently

low—while remaining within the limits imposed by the communicator's preferences and abilities.

The cognitive principle of relevance and the concept of optimal relevance on which it relies determine a comprehension procedure. Since the communicator is expected to deliver a message that should minimise the processing effort of his addressee, the addressee is entitled to follow a *path of least effort* in interpreting the communicator's message and should therefore consider cognitive effects in their order of accessibility. As soon as cognitive effects satisfy expectations of relevance, the procedure stops.

3 Relevance effects and input information

Typically, reasoning experiments involve a situation of communication: the experimenter communicates the givens of a problem (the premises, the context surrounding the premises, the instructions, the examples...) to a participant, who then communicates an answer in order to inform the experimenter about her inferential skills. Hence, the participant expects that all of what is communicated by the experimenter is relevant for reaching the answer to the provided task. Given that communicated information is expected to be relevant in solving the task, it is rational to use it for performing the task. It would indeed be foolish to presume that a piece of information provided by the experimenter is helpful for solving the task and to then ignore it. Hence, in the artificial context of a typical reasoning task, where the goal of reasoning is specified by the experimenter, all of what it is communicated should serve as input to reasoning processes. Analysing the conversational framework of the task might turn out to be helpful for identifying how participants understand it. Missing the conversational analysis can lead to overlooking that the interpretation of the task by the participant may not coincide with what it is intended by the experimenter (Henle 1962; Hilton 1995; Politzer 2004; Politzer and Macchi 2000).

Now, imagine you are an experimental psychologist and you want to design a relatively easy reasoning task that elicits a high rate of failure. How would you set about it? One simple way consists in communicating misleading information, that is, information for which expected relevance to the task is high but which actually does not help in solving it. When a piece of information is communicated by an experimenter the participant presumes that it is relevant to solve the task. The participant is then likely to end up working out a solution based on that piece of information. If however, what is believed to be highly relevant to solve the task is actually weakly relevant, the participant's ultimate strategy may not address, in an appropriate way, the problem the experimenter devised and may thus lead the participant to provide an ''erroneous'' answer. Failing to solve the task will not therefore demonstrate that the strategy itself is irrational. It will only prove that following a rational interpretation of the givens communicated by the experimenter can result in solutions that are non-optimal from the experimenter's point of view. Do tasks

containing misleading information actually exist in reasoning experiments? Yes. A good example of such a task is the 2-4-6 problem created by one of the founders of psychology of reasoning, the British psychologist Peter Wason.

Originally, Wason (1960, 1968) conceived this problem in order to investigate hypothesis-testing and to see whether lay people test hypotheses in the same way as scientists do (see Gorman 1995; Poletiek 2001, for reviews). The experimenter tells participants that he has in mind a rule that enables the generation of triples of numbers. The task the participant has to perform is to discover the rule by proposing triples to the experimenter. For each tested triple, the experimenter indicates whether or not it is compatible with the rule. When the participant thinks she has discovered the rule, she announces it to the experimenter. If the rule is incorrect, the participant is asked to continue testing triples and so on until the correct rule is provided. Before the participant starts testing triples the experimenter tells her that the triple 2-4-6 is an example of a triple generated by the rule. The rule that participants are required to discover is simply *three increasing numbers*.

Experimental findings show a low level of ''correct'' performance: In Wason's original study, only 21% of participants were able to discover the rule on their first try (Wason 1960). Typically, the rules proposed are more specific than the rule to be discovered: *three numbers increasing by 2, three increasing even numbers, three numbers increasing with equal intervals.* Participants typically test triples compatible with these hypotheses and tend to overlook triples incompatible with such hypotheses. The upshot is that participants' behaviour seems to depart from ideal scientific reasoning. To account for the low level of performance obtained in the task, researchers have often described the strategy followed by participants in terms of *cognitive biases*. According to Wason, participants exhibit a *confirmation bias*: they are unable to envisage the refutation of their hypothesis and therefore always search to confirm it. In contrast, Evans (1989) claims that participants can actually think in terms of refutation but are only able to consider positive triples (triples which are consistent with the hypothesis). A positive triple can falsify the hypothesis to the extent that the target rule is more specific than the hypothesis or that the target rule intersects with the hypothesis (see also Klayman and Ha 1987). According to Evans, the positive testing strategy is a cue revealing that participants succumb to a *positivity bias*. Poletiek (2001) favours an explanation in terms of *restrictiveness bias* and argues that positive testing results in focussing on over-restricted hypotheses. Considering such descriptions, people could be found *guilty of irrationality*. For instance, Wason (see p. 172 in 1968) depicted people who fail the 2-4-6 problem as ''narrow minded'' and ''cognitively prejudiced''.[1] However, a conversational analysis of

[1] Wason (1968) ends his chapter with the following remark: ''In the real world, as opposed to the psychological laboratory, the fixated, obsessional behaviour of some the subjects would be analogous to that of a person who is thinking within a closed system—a system which defies refutation, e.g., existentialism and the majority of religions. These experiments demonstrate, on a miniature scale, how dogmatic thinking and the refusal to entertain the possibility of alternatives can easily result in error'' (see p. 174 in Wason 1968).

the task can reveal the experimenter is actually *guilty of deception*. Below I describe why I make this accusation.

The salient properties conveyed by the example 2-4-6 strongly suggest a more specific rule than the target rule. Of course this was done deliberately in order to see whether participants were able to falsify the hypotheses suggested by 2-4-6 and to consider more general hypotheses (Wason 1968). But what makes the task deceptive is not the specificity of the triple but the fact that it is *communicated*. Consequently, it is expected to be optimally relevant and is probably viewed as an aid given by a supposedly trustworthy person (i.e., a professional psychologist who knows the rule). Given the presumption of relevance, processing the example provided by the experimenter should result in cognitive effects in the context of the task. Since the task consists in discovering a rule, any rule-like property that easily comes to mind when processing the initial triple should be considered by the participant as one the experimenter wanted him to consider in order to discover the rule. By *communicating* the triple 2-4-6, the experimenter constrains the input that will serve for working out a strategy. Features such as ''increasing by 2'' and ''evenness'' are more accessible and should therefore be considered as more appropriate for discovering the rule than the general and banal feature *increases*. If the experimenter wanted the participant not to use these features, he would have communicated a different triple. In other words, it is perfectly rational from a communicative point of view to focus on hypotheses that envelope the salient features of 2-4-6. This does mean that the target rule should directly match one of these salient properties, but it does imply that the participant's provided rule ought to exploit these properties. It then becomes extremely difficult to think about a rule that will not exhibit the salient properties of 2-4-6. Providing the triple 2-4-6 with a presumption of optimal relevance misleads participants in their expectations of relevance. Although limits in a participant's cognitive skills are likely to play a role in any reasoning task, not enough is made of the rational uses of communication that undermine ''correct'' performance. Communication has rarely been mentioned as a factor susceptible of impairing performance in the 2-4-6 task (but see Politzer 1986; Rossi et al. 2001) and its role has not been assessed empirically.

In order to determine the extent to which cognitive biases arise in the 2-4-6 task, one should neutralize the impact of its communication. That is, it would be important to devise a version of the task in which 2-4-6 is not presented with a presumption of optimal relevance. This was the purpose of an experiment I devised with Sandrine Rossi and Walter Schroyens. In one condition, the triple 2-4-6 was provided to the participant but was not communicated. We called it the ''jackpot'' condition because, before receiving any instructions about the rule discovery task, participants were faced with a ''jackpot''—a computer that generated sets of three numbers at random. In fact, in order to ascertain that participants believed the triples were randomly generated, they had to see the results of such a jackpot five times. After this introductory jackpot session, the experimenter provides the standard instructions of the rule discovery task before asking the participant to prompt the jackpot one

last time. The participant was informed that she will be told whether or not this randomly generated triple is an instance of the rule. The triple which came up conformed to the rule and was actually 2-4-6. Obviously the participant ignored that for this last trial the jackpot was biased and she was expected to believe that the regularity of the triple 2-4-6 was accidental. Since no presumption of optimal relevance accompanies this triple, its most salient properties can no longer be considered as the ones that must necessarily be taken into account for discovering the rule. Participants should therefore focus less on specific hypotheses and should explore a greater variety of triples than in a standard condition where the 2-4-6 triple is intentionally chosen by the experimenter.

The results confirmed these predictions. Correct rules were announced at higher rates in the jackpot condition than in the standard condition. In the standard condition 24% of participants proposed the correct rule on their first try whereas 55% did in the jackpot condition ($\chi^2(1) = 5.84$, $P < 0.02$). Moreover, the mean number of rules proposed before reaching the correct rule was higher in the standard condition than it was in the jackpot condition (2.38 vs. 1.59; *Mann–Whitney* $U_{29,24} = 214$, $Z = 2.39$, $P < 0.01$; see also Fig. 2) and all five participants who ultimately failed to discover the rule belonged to the standard condition. Finally, the number and the type of triples tested before proposing a rule differed across conditions. In the jackpot condition, participants proposed more triples per rule attempted than in the standard condition (8.15 vs. 6.11; *Mann–Whitney* $U_{29,24} = 202$, $Z = -2.6$, $P < 0.009$). Participants also tested a significantly higher proportion of counter-examples (i.e., triples that do not fit to the target rule) and of irregular increasing triples (i.e., triples which do not increase with the same interval) in the jackpot condition than in the standard condition (see Van der Henst et al. 2002a for details). These data show that participants of the jackpot condition examine a greater variety of hypotheses than participants in the standard condition, who tend to test hypotheses fitting with the specific features conveyed by 2-4-6.

The results confirm the conversational hypothesis based on relevance. It cannot be argued that the relatively weak performance commonly observed in the standard task is due uniquely to cognitive biases. A substantial proportion

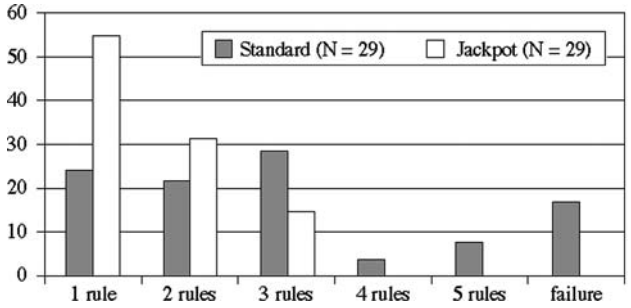

Fig. 2 Number of rules announced before reaching the correct solution in the Standard and Jackpot conditions

of non-optimal responding is due to the experimenter's misleading communication. By presenting the task in a less misleading way, that is, by removing the presumption of optimal relevance in the jackpot condition, participants perform more ''rationally''.

In a second experiment we followed a different strategy: instead of diminishing the expected relevance of the salient features of 2-4-6, we enhanced the relevance of the property *increase* (Van der Henst et al. 2002a). We embedded the rule discovery task in the context of economic activity. For an economic agent, ''increase'' is the key word; an economic agent always wants to increase sales, profits, productivity and so on. In the context we used, participants had to discover, not an artificial rule an experimenter has in mind but, a rule about car sales that a dealer wants to enforce (see Fig. 3). We presumed that in such contexts, participants will be less likely to focus on hypotheses that display the salient properties of 2-4-6 than in the control condition. The procedure differed from previous experiments since participants did not interact with an experimenter but with a computer. They had to enter triples by typing them with the keyboard. Triples appeared in green

Mister Jansens is a prosperous garage owner. He has recently posted an advertisement in order to recruit a car salesman. Bert answered the ad and obtained an interview with Mr Jansens; Bert told him that he was very motivated for the job but he also informed Mr Jansens that he never sold cars before. Mr Jansens has a good feeling about Bert, but he thinks that an interview is not good enough for deciding if Bert will be a good car salesman. Hence, Mr Jansens offers Bert a job in his agency for three months as a sales-training period. If Bert does a good job, then he will decide to hire him. In particular, in order to be recruited, Bert's sales during these three months are required to follow a rule that is important to Mr Jansens. Bert willingly accepts the proposition.

Three months later… It's time to check out how Bet did! The first month, Bert was able to sell 2 cars, the second month 4 cars and the third 6 cars. The verdict of Mr Jansens is very clear: "Perfect! Your sales have well respected the rule. So, I can hire you!

Fig. 3 The context of economic acitvity

when they were consistent with the rule and in red when inconsistent. Participants could propose only one rule.

More participants discovered the target rule in the economic condition than in the control condition (29.3% vs. 3.2% $\chi^2(1) = 15.3, P < 0.001$). Additionally, participants tested more triples and more variety of triples in the economic condition than in the control one (see Van der Henst et al. 2002a for details). What is however surprising in this experiment is the low level of performance obtained in the two conditions. Wason's initial study and other subsequent studies report a much higher level of success with the standard abstract task (10–30% of participants gave the correct rule on their first try). This discrepancy may result from the experimental procedure we introduced. Given that participants were not face-to-face with an experimenter, but simply interacted with a computer, they could act in more anonymous way and might therefore get less involved in the task. In particular, they might feel less embarrassed in the case they provide an incorrect answer and might be more likely to adopt a risky strategy by proposing a rule after only a few triples were tested.

To conclude, people tailor their search of hypotheses according to what they consider to be relevant. When 2-4-6 is presented as optimally relevant for dealing with the standard task, participants focus on specific hypotheses and tend to fail the experimenter's intended task. When the presumption of optimal relevance that accompanies the triple 2-4-6 is removed, or when the concept of *increase* becomes more relevant, participants focus less on specific hypotheses linked to the salient properties of 2-4-6 and succeed in providing ''correct'' responses on the task. This does not mean that cognitive biases do not exist, but in the absence of conversational analysis, these biases can be easily overestimated. This study is in line with research analysing the role of communication in other reasoning tasks (e.g., see Dulany and Hilton 1991; Politzer and Noveck 1991, on the Linda problem; Schwartz et al. 1991; Macchi 1995, on the base-rate fallacy; Sperber et al. 1995 on the Wason selection task).

4 Relevance effects and output information

What conclusions do people intend to attain when they reason? What makes a conclusion worth being inferred? For sure, a reasoner aims at deriving sound conclusions. The soundest conclusions a reasoner can infer are conclusions that logically follow from the premises: if the premises are true, a conclusion that logically follows from the premises will necessarily be true. One could therefore argue that people aim at deriving only logical conclusions. However, such a criterion seems to be both too strict and too loose. First, many inferred conclusions are not logically valid but are perfectly justified. Take for instance conversational implicatures (Grice 1975). These inferences are pragmatically justified and the audience must even derive them in order to capture what the speaker means. Imagine that I tell you that *some* of my friends are married. You will probably infer that *not all* of my friends are married. Such a conclusion is not logically valid since from a logical point of view ''some'' could be

compatible with ''all''. However, from a pragmatic point of view, if I choose to use the weaker informative term some rather than the more informative term *all*, this means than all was not appropriate and justifies the *not all* inference. It is therefore not necessary for a conclusion to be logical in order to be rationally derived.

Second, an infinity of conclusions follow from a given set of premises and most of these conclusions are pointless. From the single premise P, no one will aim at drawing the following logical conclusions:

P & P
(Not (Not (Not (Not (Not (Not P))))))
P or Q or S or T or U or V

Deriving such conclusions would be cognitively inefficient because cognitive resources would then be allocated to worthless information. When we reason, we are not interested in a conclusion only because it is sound but mainly because it matters for us. Harman (1995) has expressed this idea with what he called the *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 (p. 186).

It is therefore not sufficient for a conclusion to be logically valid in order to be worth being inferred. Inferring a conclusion does not only depend on our cognitive abilities but also on the cognitive purposes we pursue. In two recent papers, Van der Henst et al. (2002b) and Van der Henst and Sperber (2004) proposed that, in line with the first principle of relevance, people tend to draw relevant conclusions. In the present paper, I will present four unpublished experiments corroborating this view.

Imagine that you hear the following conversation:

– If John Hunter was faster than Bob Clark in this race then, the new 100 m Olympic champion has a first name beginning with J. (if p then q)
– It turns out that John Hunter was faster than Bob Clark in this race. (p)

What conclusion will you infer? A very compelling conclusion is the one that results from a Modus Ponens inference:

[4] The new 100 m Olympic champion has a first name beginning with J.

Such a conclusion is logically valid and it is perfectly within the scope of basic human inferential abilities: as described in many conditional reasoning studies, from 90 to 100% of people easily perform the Modus Ponens inference. But what about the relevance of that conclusion?

In a reasoning experiment, participants are expected to pretend that the premises and some conclusions might be relevant for them. Let's pretend that you have some interest in attaining knowledge about the new Olympic

champion of 100 m. In such ordinary contexts, knowing that the new of Olympic champion of 100 m has a first name beginning with a J is likely to be weakly relevant. For sure, it will allow you to infer that many other athletes cannot be the Olympic champion (those whose first name does not start with a J) but it will not provide you with the name of the champion itself, a piece of information you may well be determined to know in such contexts. Thus, you probably would go beyond this weakly relevant conclusion and infer that:

[5] The new Olympic champion is John Hunter.

Although this conclusion is not logically valid it is pragmatically justified. When one says *if P then Q*, one usually means that believing in *P* is good enough for believing *Q*. Let's try to figure out how the fact that John Hunter was faster than Bob Clark in a certain race could be a good reason to believe that the new 100 m Olympic champion has a first name beginning with J. First, by the very fact that the race provides definite information about the name of the new 100 m Olympic champion, one may reasonably assume that this race is the Olympic 100 m final. Second, because the greater speed of John Hunter over Bob Clark enables the belief that the new Olympic champion has a first name beginning with J, and because John Hunter has a first name starting with a J, one may reasonably assume that John Hunter and Bob Clark were the two leaders of the race and that John Hunter arrived first. These additional assumptions are easily accessible and give sense to the conditional statement. This endows them with a relatively high degree of plausibility. In fact, other assumptions can account for the link between the antecedent and the consequent of the conditional statement but are less immediately accessible and are thus less plausible. For instance, it may be the case that some other athletes whose first name also starts with a J arrived before John Hunter at the Olympic final. In short, in order to make sense of the conditional statement, one may interpret it as meaning *If John Hunter was faster than Bob Clark in this race (i.e., the Olympic 100 m final), then the new Olympic champion is John Hunter*. From this conditional and the second premise, conclusion [5] follows. This conclusion is more relevant than conclusion [4] because it conveys more cognitive effects. Indeed, any cognitive effect that can be achieved with [4] can also be achieved with [5] but most cognitive effects that can be achieved with [5] cannot be achieved with [4].

What do people, who participate in a reasoning experiment, do? Do they play the game of logic and infer a weakly relevant but logical conclusion? Or do they play the game of relevance and infer a more relevant but only pragmatically justified conclusion? In a first experiment, 55 participants were given the two above-mentioned premises and were asked to answer the following question ''What can you deduce from this information?'' Figure 4 shows the percentages of participants for each category of conclusions. It appears clear that more participants aimed at inferring a relevant and pragmatically justified conclusion than a logical and weakly relevant conclusion. Interestingly, some participants were aware that the pragmatic conclusion was

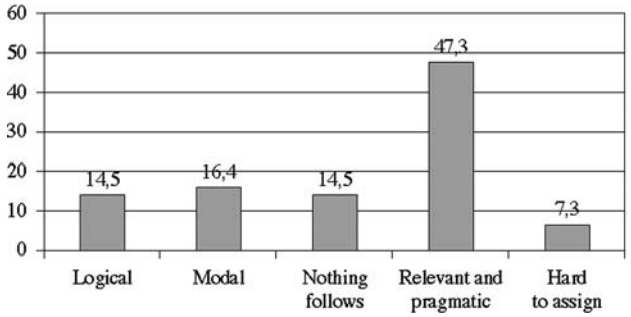

Fig. 4 Percentage of conclusion types

less guaranteed than the logical one. They were reluctant to express the unsound conclusion [5] or the irrelevant conclusion [4] and choose to produce a modal conclusion (e.g., ''John Hunter, given that his first name starts with a J, could be the 100 m Olympic champion'', ''Bob Clark lost, we can assume that John Hunter is the 100 m Olympic champion'') or a ''nothing follows'' answer (''Nothing, there might have other racers whose first name starts with a J. Is the race the Olympic race? We don't know'', ''Nothing, no conclusion. It means in no way that the new 100 m Olympic champion could be John Hunter''). Modal conclusions have the advantage of being more relevant than [4] and sounder than [5].

In a second experiment, a control group of 60 participants received the same premises and the same question as in the first experiment while an experimental group of 61 participants received the same premises with the following question: ''what can you deduce with complete certainty?'' Such a question specifies more clearly that the participant's conclusion can achieve relevance by revealing to the experimenter her ability to perform a strictly logical inference. As expected, significantly more participants produced a logical conclusion in the experimental group than in the control group (37.7% vs. 13.3%, $\chi^2(1) = 9.43$, $P < 0.005$; see Fig. 5). Interestingly, in the experimental group, a significant proportion of logical answers were based on a single premise (''John Hunter ran faster than Bob'', ''John Hunter arrived before Bob Clark'', ''One can only say that Bob Clark arrived after John Hunter''). Although such a conclusion is a simple repetition or a restatement

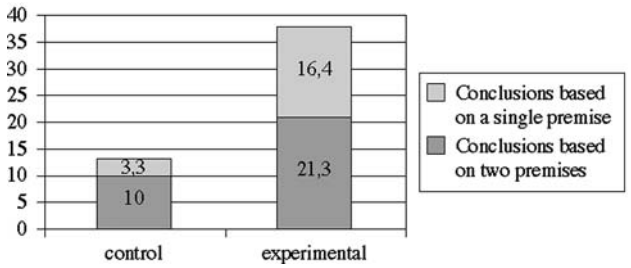

Fig. 5 Percentages of logical conclusions in the control and experimental conditions

of one premise, participants who formulated it may have estimated that it was more relevant than the Modus Ponens conclusion.

A third experiment aimed at manipulating the rates of logical conclusion by manipulating the amount of cognitive effects it could convey for the reasoner. Two problems for which the Modus Ponens conclusion differed in terms of cognitive effects were given to two different groups of participants:

Problem A

- If Cameroon has beaten Brazil in this match then the country who won the football world cup has a name beginning with a C.
- Cameroon has beaten Brazil in this match.

Problem B

- If Cameroon has beaten Brazil in this match then the country who won the football world cup is an African country.
- Cameroon has beaten Brazil in this match.

The Modus Ponens inference following the two premises of Problem A and Problem B are respectively:

[6] The country who won the football world cup has a name beginning with a C.
[7] The country who won the football world cup is an African country.

For both problems, the pragmatically justified conclusion - *the country who won the football world cup is Cameroon*—is more relevant than the Modus Ponens conclusion since it conveys more cognitive effects. However, conclusion [7] is more likely to address typical concerns one may have about the winner of the football world cup than conclusion [6]. Conclusion [7] should therefore more often be produced than conclusion [6]. This is indeed what was observed: participants who received Problem A formulated the Modus Ponens conclusion significantly more often than those who received Problem B (22.3% vs. 34% $\chi^2(1) = 5.14$, $P < 0.03$, $N = 301$).

5 Conclusion

The experiments reported in this article show that the selection of input information and the production of output information in reasoning give rise to relevance effects. First, an experimenter can easily mislead participants and prevent participants' success by indicating that some information is relevant to solving a reasoning task when, in fact, it is not. Data obtained on the 2-4-6 task show that, when one weakens the expected relevance of unhelpful information or when one increases the relevance of helpful information, performance increases. These data are in line with recent studies pointing to the importance pragmatic factors in solving reasoning problems (see Hilton 1995; Politzer 2004; Politzer and Macchi 2000 for reviews). The pragmatic analysis can partially explain why performance in relatively easy task can be so low. An-

other example that illustrates this claim comes from the *Linda* problem devised by Tversky and Kahneman (1983). This task elicits a high level of failure while it is computationally simple. Indeed, to solve it, one only needs to recognise that the probability of two events ('A' & 'B') cannot be higher than the probability of one of these events (A): $P(A \& B) \leq P(A)$. However, the task is presented in such way that event (B) is likely to be considered as highly relevant for solving the task so that participants tend to interpret the single event (A) as the conjunction of two events (A & not-B). When this interpretation is discouraged, rates of correct performance increase (see Politzer and Noveck 1991). However, the manipulation of pragmatic factors in those tasks does not imply that the experimenter can act at will on the level of performance. As indicated by the first experiment, although the jackpot condition increased correct performance, about one out of two participants still announced an incorrect rule on their first try. Cognitive biases may be attenuated but not necessarily removed.

Second, in the conditional study reported above, relevance appears to be a key factor in the derivation of conclusions. Participants tend to reach the most relevant conclusions for them. A logical but weakly relevant conclusion is less likely to be drawn than a pragmatically justified and highly relevant conclusion. The prominence of relevance in reasoning does not imply that relevance is the only factor that affects the derivation of conclusions. As pointed out in Introduction, the nature of the processes involved in reasoning also determines the conclusions that can be drawn. Considerations of relevance are however helpful in specifying the restriction that reasoners put on the conclusion they actually produce. This restriction issue is generally ignored in the field of psychology of reasoning. However, in their seminal book on deduction Johnson-Laird and Byrne (1991) are among the rare researchers to have addressed it in the past. They judiciously argue that "when reasoners make a deduction in daily life, they must be guided by more than logic" (p. 21). According to them three extra-logical constraints govern reasoning: conclusions should not throw semantics away, they should be parsimonious and they should assert something new. However, these constraints may not be sufficient to predict the conclusions produced in the conditional study reported here. The Modus Ponens conclusion following from the "race problem" above was rarely stated by participants; although such a conclusion satisfies the three constraints, people stated a more relevant conclusion. In short, relevance might be a better candidate for being this "more than logic" component.

Acknowledgments The author would like to thank Ira Noveck and an anonymous reviewer for helpful comments on a previous version of this paper.

References

Braine MDS, O'Brien DP (1998) Mental logic. Lawrence Erlbaum Associates, Mahwah
Clark HH (1969) Linguistic processes in deductive reasoning. Psychol Rev 76:387–404
Dulany DE, Hilton DJ (1991) Conversational implicature, conscious representation, and the conjunction fallacy. Soc Cogn 9:85–110

Evans JStBT (1989) Bias in human reasoning: causes and consequences. Lawrence Erlbaum Associates, Hove
Evans JStBT (1995) Relevance and reasoning. In: Newstead SE, Evans JStBT (eds) Perspectives on thinking and reasoning. Essays in honour of PeterWason. Lawrence Erlbaum Associates, Hove
Gorman ME (1995) Hypothesis testing. In: Newstead SE, Evans JStBT (eds) Perspectives on thinking and reasoning. Essays in honour of Peter Wason. Lawrence Erlbaum Associates, Hove
Grice HP (1975) Logic and conversation. In: Cole P, Morgan JL (eds) Studies in syntax, vol 3: Speech acts. Academic, New York
Hardman D (1998) Does reasoning occur on the selection task? A comparison of relevance-based theories. Think Reasoning 4(4):353–376
Harman G (1995) Rationality. In: Smith EE, Osherson DN (eds) Thinking: an invitation to cognitive science, vol 3, 2nd edn
Henle M (1962) The relation between logic and thinking. Psychol Rev 69:366–378
Hilton DJ (1995) The social context of reasoning: conversational inference and rational judgment. Psychol Bull 118:248–271
Huttenlocher J (1968) Constructing spatial images: a strategy in reasoning. Psychol Rev 75:550–560
Johnson-Laird PN (1983) Mental models. Cambridge University Press, Cambridge
Johnson-Laird PN, Byrne RMJ (1991) Deduction. Lawrence Erlbaum Associates, Hove
Johnson-Laird PN, Byrne RMJ (2002) Conditionals: a theory of meaning, pragmatics, and inference. Psychol Rev 109:646–678
Johnson-Laird PN, Savary F (1999) Illusory inferences: a novel class of erroneous deductions. Cognition 71:191–229
Klayman J, Ha YW (1987) Confirmation, disconfirmation and information in hypothesis testing. Psychol Rev 94:211–228
Macchi L (1995) Pragmatic aspects of the base-rate fallacy. Q J Exp Psychol 48:188–207
Poletiek F (2001) Hypothesis-testing behavior. Psychology Press—Taylor and Francis, Philadelphia
Politzer G (1986) Laws of language use and formal logic. J Psycholinguist Res 15:47–92
Politzer G (2004) Reasoning, judgment, and pragmatics. In: Noveck I, Sperber D (eds) Experimental pragmatics. Palgrave Macmillan, Basingstoke
Politzer G, Macchi L (2000) Reasoning and pragmatics. Mind Soc 1:73–93
Politzer G, Noveck IA (1991) Are conjunction rule violations the result of conversational rule violations. J Psycholinguist Res 20:83–103
Rips LJ (1994) The psychology of proof. The MIT, Cambridge
Rossi S, Caverni JP, Girotto V (2001) Hypothesis testing in a rule discovery problem: when a focused procedure is effective. Q J Exp Psychol A 54:263–267
Schwartz N, Strack F, Hilton D, Naderer G (1991) Base rates, representativeness, and the logic of conversation: the contextual relevance of ''irrelevant'' information. Soc Cogn 9:67–84
Sperber D, Wilson D (1995) Relevance: communication and cognition. Blackwell, Oxford
Sperber D, Cara F, Girotto V (1995) Relevance theory explains the selection task. Cognition 57:31–95
Tversky A, Kahneman D (1983) Extensional vs. intuitive reasoning the conjunction fallacy in probability judgment. Psychol Rev 91:293–315
Van der Henst JB, Sperber D. (2004) Testing the cognitive and communicative principles of relevance. In: Noveck I, Sperber D (eds) Experimental pragmatics. Palgrave Macmillan, Basingstoke
Van der Henst JB, Rossi S, Schroyens W (2002a) When participants are not misled they are not so bad after all: a pragmatic analysis of a rule discovery task. Proceedings of the 24th annual conference of the cognitive science society
Van der Henst JB, Sperber D, Politzer G (2002b) When is a conclusion worth deriving? A relevance-based analysis of indeterminate relational problems. Think Reasoning 8:1–20
Wason PC (1960) On the failure to eliminate hypotheses in a conceptual task. Q J Exp Psychol 12:129–140
Wason PC (1968) On the failure to eliminate hypotheses. A second look. In: Wason PC, Johnson-Laird PN (eds) Thinking and reasoning. Penguin, Baltimore
Wilson D, Sperber D (2004) Relevance theory. In: Horn LR, Ward G (eds) The handbook of pragmatics. Blackwell, Oxford. pp 607–632