

Uncertainty and the suppression of inferences

Guy Politzer

► **To cite this version:**

Guy Politzer. Uncertainty and the suppression of inferences. Thinking and Reasoning, Taylor & Francis (Routledge), 2005, 11 (1), pp.5-33. <ijn_00000616>

HAL Id: ijn_00000616

https://jeannicod.ccsd.cnrs.fr/ijn_00000616

Submitted on 30 Jul 2005

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.

Thinking and Reasoning,
2005, 11, 5-33

Uncertainty and the Suppression of Inferences

Guy Politzer

Centre National de la Recherche Scientifique (CNRS)

Saint-Denis, France

Abstract

The explanation of the suppression of Modus Ponens inferences within the framework of linguistic pragmatics and of plausible reasoning (i. e. , deduction from uncertain premises) is defended. First, this approach is expounded, and then it is shown that the results of the first experiment of Byrne, Espino and Santamaría (1999) support the uncertainty explanation but fail to support their counterexample explanation. Second, two experiments are presented. In the first one, aimed to refute one objection regarding the conclusions observed, the additional conditional premise (*if N, C*) was replaced with a statement of uncertainty (*it is not certain that N*); the answers produced by the participants remained qualitatively and quantitatively similar in both conditions. In the second experiment, a fine-grained analysis of the responses and justifications to an evaluation task was performed. The results of both experiments strongly supported the uncertainty explanation.

Uncertainty and the Suppression of Inferences

Because their sources of information generally are not fully reliable, people are more or less confident in the truth of the propositions that they entertain. Consequently, an important part of everyday reasoning, possibly the major part, is based on uncertain premises. This applies, in particular, to deduction, with the important consequence that people draw uncertain conclusions. Indeed, although deductive arguments are characterised by the delivery of a sure conclusion (if the premises are true), a deductive argument that has one or more uncertain premises delivers an uncertain conclusion.

Psychological investigations of plausible reasoning

The psychological investigation of plausible reasoning, that is, deduction from uncertain premises, is relatively recent, in spite of the pervasiveness of this kind of inference, and of the interest shown by logicians and mathematicians in the past (e. g., Boole, 1854/2003) and more recently (Adams, 1998; Adams & Levine, 1975; Nilsson, 1986; Rescher, 1976; Suppes, 1966). Even though the psychological studies of plausible reasoning have been limited in scope, they have produced at least one clear and consistent result: When the conditional premise of a simple argument such as Modus Ponens (MP) or Modus Tollens (MT) is uncertain, the uncertainty propagates to the conclusion. The most significant of these studies will be reviewed very shortly.

George (1995) had the credibility of causal conditional statements (such as *if the winter is harsh, then many people will be ill*) evaluated on a confidence scale; then, the same statements were used as premises of arguments whose conclusion had to be evaluated on the same scale. A high correlation between the two evaluations was observed. In another experiment, the participants who were invited by the experimenter to consider that the conditional premise was true endorsed the conclusion significantly more often than those who were invited to consider this premise as debatable. Thompson (1994, 1995) observed that the rate of endorsement was an increasing function of the degree of “perceived sufficiency” of the conditional premise independently estimated by judges. Liu, Lo, & Wu (1996) arrived at the same conclusion in using the conditional probability of the premise’s consequent on its antecedent as an operational

definition of the “perceived sufficiency”. Cummins (1995; Cummins, Lubart, Alksnis, & Rist, 1991) observed that, for causal conditionals, as the number of disabling conditions increases, the level of confidence in the conclusion decreased (disabling conditions are events whose realization would prevent the effect from occurring; they may be regarded as responsible for the lack of perceived sufficiency of the causal conditional). Stevenson & Over (2001) manipulated the confidence in the conditional premise by using contrasted scenarios in which the premise was uttered by an expert or by a novice: The perceived likelihood of the conclusion was greater in the former case than in the latter. George (1997) used arguments whose conditional premise was formulated with probability expressions such as *it is very likely that* or *it is not very likely that*, etc. , and asked participants to evaluate the conclusion on a scale ranging from *certainly false* to *certainly true*; nearly all the answers expressed uncertainty and the modal answer was identical to the probability expression used in the premise. In brief, in all these studies, various degrees of belief in the conditional premise of simple arguments like MP and MT were defined by an independent evaluation or by manipulation; this resulted in a degree of confidence in the conclusion that closely reproduced that of the premise, and in rates of endorsement of the conclusion that were an increasing function of the degree of belief in the premise.

The suppression effect

The studies just reviewed, which all involved conditional premises, followed the revival of interest in the study of reasoning with meaningful conditionals that took place in the late 1980's (e. g. , Byrne, 1989; Byrnes & Overton, 1986; Clement & Falmagne, 1986; Dias & Harris, 1988; Markovits & Vachon, 1989; O'Brien, Costa & Overton, 1986). Byrne (1989) presented a new paradigm using modified MP and MT arguments. One control group solved standard arguments, such as, for MP: *If she has an essay to write, then she will study late in the library; she has an essay to write. Conclusion: (a) she will study late in the library; (b) she will not study late in the library; (c) she may or may not study late in the library.* Another group solved the same argument modified by the addition of a second conditional premise (the *additional premise*) chosen in such a way that its

consequent was the consequent of the major premise, and its antecedent was a *necessary* condition for the consequent to occur, such as *If the library stays open, then she will study late in the library*. Whereas for the control group the rate of correct answers (a) was above 95%, as usually observed, for the other group, less than 40% of the participants chose this answer. Thus, it would seem that the valid inference of MP can be suppressed; a similar effect was observed for MT. This suppression effect has been replicated many times (Bonneton & Hilton, 2002; Chan & Chua, 1994; Dieussaert, Schaeken, Schroyens, & d'Ydewalle, 2000; Stevenson & Over, 1995).

The uncertainty account of the suppression effect.

The pragmatic mechanism. In a theoretical discussion of Byrne's (1989) claim that the valid argument of MP can be suppressed, and of one of its consequences, viz., there might be no mental rule for MP, Politzer and Braine (1991) outlined an explanation of the suppression effect on which the present paper will be focused (note 1). This explanation hinges on a pragmatic analysis of the conditional premises. Let *if A, C* and *if N, C* stand for the major premise and the additional premise, respectively (*N* being a necessary condition for *C* to occur). Politzer and Braine remarked that ordinary conditionals such as the major premise *if A, C* are typically uttered with a *ceteris paribus* assumption to the effect that the set of conditions necessary for *C* to hold are satisfied. This assumption can be characterized as an *epistemic implicature* to the effect that the speaker does not believe or does not know that these necessary conditions do not hold. For instance, it would be deceptive for the speaker to state *If she has an essay to write, then she will stay late in the library*, while knowing that the library is not open (and similarly for many other conditions that would prevent the person from studying in the library). Also, upon being told that *if N, C*, the hearer (or for that matter, the participant) may exploit another property of conditional sentences: In an appropriate context, they license the generation of an implicature - again epistemic - to the effect that the speaker does not know whether or not the antecedent (*N*) is satisfied. (In case the speaker knew it was, she should have used the term *since* instead of *if*, or utter the consequent unconditionally).

These two properties, which concern the major premise and the additional premise, respectively, interact when these premises are brought together: A doubt about the satisfaction of N is communicated by the additional premise: This cancels the implicit assumption of the satisfaction of N which accompanies the major premise, and so the truth of the major premise is no longer warranted. Whereas in the absence of the additional premise (in the standard MP) the conclusion C follows, in the presence of the additional premise the major premise is uncertain and this uncertainty propagates to the conclusion, which participants are reluctant to draw. Of course, this process is not deterministic: As for any implicature, the generation of the epistemic implicature attached to the additional premise is under the hearer's/participant's responsibility: Some may not generate it and treat the argument as a standard MP, drawing the conclusion C . In brief, once the source of the uncertainty of the conditional premise has been identified, the suppression task can be subsumed under plausible reasoning tasks.

Empirical support. This explanation of the suppression effect has received considerable support, indirect and direct. Indirect support can be found in Chan and Chua's (1994) results. They used different necessary conditions that varied in importance (estimated by judges), and observed that the rate of endorsement of the conclusion was a decreasing function of the importance of the necessary condition: By manipulating levels of necessity, they introduced levels of belief in the major premise.

Neth and Beller (1999) used pairs of statements (A and C) that had been estimated by judges to be in a relation of sufficiency, but not of necessity. Participants were asked to decide which of the following sentences was appropriate to express the relation between A and C : (a) *If A then C* ; (b) *If C then A* ; (c) both of these; (d) none of these. In a control group, nearly all the participants chose option (a); but in a group in which an additional statement said that no information was available about a condition N necessary for C to occur, 40% of the participants chose option (d), which demonstrates the loss in assertability of the conditional *if A , C* that results from the uncertainty about the satisfaction of N .

More direct support was obtained by Stevenson and Over (1995) who replaced the additional premise *if N, C* with a categorical premise that had various levels of certainty: *Always N, almost always N, sometimes N*, etc. They observed that, as the frequency level of the additional premise decreased, the rate of endorsement of the conclusion decreased, and participants' uncertainty scores increased monotonically. They concluded, in line with Politzer and Braine's (1991) claim, that the role of the additional premise in the suppression paradigm is to introduce uncertainty in the major premise.

Manktelow and Fairley (2000) using deontic conditional rules replaced the additional premise with a categorical premise that had various levels of satisfaction for *N* (rather than frequency terms). Again, the degree of certainty of the conclusion of MP arguments was an increasing function of the degree to which *N* was satisfied.

Politzer and Bourmaud (2002) using various conditional rules (causal, means-end, remedial, decision) generalized the foregoing results. They observed that the rate of endorsement of the conclusion and the degree of belief in the conclusion of MT arguments were an increasing function of the degree of satisfaction of the necessary condition and of its importance. Even more directly, Bonnefon and Hilton (2002) manipulated the probability for the additional premise to cause doubt about the necessary condition. They did so by varying the syntactic form of the conditional premise. Whereas the participants who were presented with either one of the additional premises *if N, C* (the usual one), *if not-N, not-C*, and *if not-C, not-N*, had a lower degree of belief in the conclusion than had the participants who received a standard MP, no difference was observed between the latter participants and those who were presented with the additional premise *if C, N*. Now, the linguistic analysis, confirmed by participants' judgments collected after completion of the inference, indicates that, contrary to the latter additional premise, the first three arouse doubt about *N* (the first one by generating an epistemic implicature, the other two by explicitly mentioning *not-N*).

In summary, there is a very strong case for the notion that the suppression effect is an instance of deductive reasoning from uncertain premises: When a

premise added to the major premise *if A, C* introduces doubt about the satisfaction of a condition (*N*) that is necessary for *C* to hold, the credibility of the major premise is questioned and the conclusion inherits a level of credibility which depends on that of *N*.

Specifying the propagation of uncertainty. There is, however, a possible objection to the foregoing explanation which, as far as MP is concerned, can be raised on theoretical grounds. Suppose a doubt is raised about *N* (e. g., a doubt that the library is open): On one hand, it is claimed that Mary's having an essay to write is no longer a sufficient reason for the hearer to believe that she will stay late in the library, that is, the credibility of the conditional has decreased, and the credibility of the conclusion decreases accordingly, as verified by all the experimental results. But on the other hand, the very fact that one doubts that the library is open directly casts doubt on the possibility for Mary to stay in the library. That is, the doubt about the conclusion of MP can be explained more directly by the doubt about *N* without resorting to the doubt about the conditional (which itself is induced by the doubt about *N*). So, we need to discriminate between the uncertainty about *N* and the uncertainty about the conditional *if A then C* as causal factors of the uncertainty about the conclusion *C*. There are several interrelated questions which we are going to examine from a formal point of view.

1. The first question is: Does a doubt about *N* imply a doubt about *C*? The answer is straightforward and it is positive: This is because *C* logically implies *N* (*N* is a necessary condition for *C* to occur) and therefore $p(C) \leq p(N)$. As $p(N)$ decreases, so must $p(C)$ for there is no intermediate fixed lower bound for $p(N)$: *Belief in the conclusion must decrease as the satisfaction of the necessary condition N is doubted.* From this result it might seem, at first sight, that the suppression effect can be explained parsimoniously by the hypothesis that as a doubt about the satisfaction of *N* is suggested by the additional conditional, the credibility of the conclusion is directly lowered. There is, however, an essential objection to the generality of this mechanism; this objection must be considered to the extent that we are looking for mechanisms applicable beyond the special case of MP. MP is special in that its conclusion, *C*, is confounded with the consequent of the major

premise. Now take the case of MT and consider what happens to the credibility of its conclusion, *not-A*, as $p(N)$ decreases. The belief in $p(\text{not-}A)$ (she does not have an essay to write) must decrease (and this was found in the experiments reviewed). But there is no *direct* relation between these two probabilities. When N is doubted (the library may not be open) this cannot directly affect the belief about *not-A* (the opening of the library does not directly affect the probability that Mary has or does not have an essay to write). But notice this essential point: As is the case with MP, the doubt about N does affect the belief in the conditional (one cannot have full belief that if she has an essay to write she will study late in the library, while knowing that the library might not be open). Can this be accounted for? This is the next question to be formally examined.

2. In what follows, we need only to assume that belief in the conditional premise is measured by the conditional probability of the consequent on the antecedent. Theoretically, this assumption is based on a long tradition in linguistics and in philosophical logic (Adams, 1965, 1966; Lewis, 1986; Ramsey, 1931/1990). Psychologically, it has been proposed by a number of investigators (Liu & al. , 1996; Oaksford & Chater, 2003; Oaksford, Chater & Larkin, 2000) and it is supported by recent experimental data (Evans, Handley & Over, 2003; Oberauer & Wilhelm, 2003). We consider the following equation and inequalities:

(i) $p(C/A) = p(A\&C) / p(A)$ (by definition of conditional probability)

(ii) $p(A\&C) \leq p(C)$ (a theorem of probability theory).

(iii) $p(C) \leq p(N)$ (repeated from above)

(ii) and (iii) together yield: $p(A\&C) \leq p(N)$. For the same reason as above, it follows that after $p(N)$ has decreased and reached a value equal to $p(A\&C)$, $p(N)$ and $p(A\&C)$ must decrease together; but since belief in the conditional is $p(A\&C)$ divided by a constant term ($p(A)$), this means that *belief in the conditional must decrease as belief in N decreases*. A more rigorous demonstration of this point can be found in Politzer & Bourmaud (2002).

The last question is: Can we formally account for the decrease in belief in the conclusion when the belief in the conditional premise is lowered? Probabilistic logics aim to answer such questions. For probabilistic MP, the answer can be given in the form of a probability interval whose bounds depend

on $p(A)$ and on the credibility of *if A, C*. In the special case that is most often considered in psychological studies, that is when $p(A) = 1$, the answer is remarkably simple: *Belief in the conclusion equals belief in the conditional* (whether the definition of belief in the conditional is the conditional probability or the probability of the truth of the associated material conditional) (note 2). The case of MT differs. If the probability of the conditional is assumed to be the probability of the truth of the associated material conditional, the probability of the contrapositive *if not-C, not-A* is equal to the probability of *if A, C*. It follows that the degree of belief in *not-A* is equal to the degree of belief in *if A, C*. But if the probability of the conditional is assumed to be the associated conditional probability, things are notoriously more difficult as belief in *if not-C, not-A* need not be equal to belief in *if A, C*: For lack of knowledge of $p(\text{not-A} \ \& \ \text{not-C})$ one cannot know the probability of the contrapositive. However, there is an important qualitative result: When *if A, C* is uncertain, so must be its contrapositive. The reason is that *if A, C* and *if not-C, not-A* have the same necessary and sufficient condition to be uncertain, namely $p(A \ \& \ \text{not-C}) > 0$. As a consequence, uncertainty about the conditional results in uncertainty about the contrapositive, which, given *not-C*, is transmitted to the conclusion *not-A*.

The answers given to the questions above are normative. It is assumed that reasoners follow the three propositions that are italicized. These should be distinguished from the conditional probability hypothesis, which normatively is a neutral choice.

To summarize: A doubt about *N* always decreases the credibility of the conditional; a decrease in the credibility of the conditional always lowers the credibility of the conclusion; but only for MP does a doubt about *N* decrease the credibility of the conclusion. We take the credibility of the conditional as the causally relevant variable that explains the variation in belief in the conclusion. However, in the case of MP, it is not theoretically or empirically feasible to distinguish what causes the decrease in the credibility of the conclusion, a direct link from *N* or an indirect link through the belief in the conditional, as these links are correlated. We now turn to another explanation of the suppression effect.

The counterexample account of the suppression effect.

Byrne, Espino and Santamaría (1999, henceforth BES) proposed a different account, based on the notion of *counterexample* within mental model theory. According to this account, because background knowledge indicates that the antecedents of the two conditionals *if A, C*, and *if N, C* refer to additional conditions for the same outcome, people form a single model of these antecedents:

A N

following which the conditionals are represented as

A N C

...

where the dots indicate that there may be alternatives to the first model, but these alternatives are not represented explicitly (this would exceed working memory capacity). People keep in mind the initial model which they may "flesh out" to represent the following alternatives:

A N C

A not-N not-C

from which they do not endorse the conclusion *C* because of the availability of a model in which *A* is the case but it is not the case that *C*, a situation called a counterexample (BES, p. 351). The construction of this model follows from the "principle of pragmatic modulation" (Johnson-Laird & Byrne, 2002): The model not-N not-C is assumed to be explicitly represented in long term memory and presumably is integrated with *A*. In brief, the "key suggestion is that the provision of an [...] additional antecedent enables people to represent counterexamples to conclusions explicitly" (BES, p. 352) (note 3).

BES presented four experiments, the first and the third of which are specially relevant to the present discussion. In the third experiment, they tested a hypothesis derived from mental model theory: A biconditional formulation of the additional premise should result in more suppression than the conditional formulation because a biconditional is semantically stronger than a conditional (it eliminates more alternative situations), and it emphasizes the importance of the necessary condition in the additional premise. This prediction was satisfied. Of

course, it is also immediately derivable from the uncertainty account, as the biconditional formulation (*if and only if N*) makes it explicit that *N* is a necessary condition, so that it suggests more strongly the epistemic implicature to the effect that *N* might not be satisfied. In sum, this experiment does not provide a test between the two approaches (neither was it designed for this purpose), as both can easily explain its results. In contrast, BES's first experiment, which used a production task, was an attempt to refute the uncertainty explanation. The main aim of the present paper (besides expounding the uncertainty account) is to demonstrate that this attempt failed: Based on the two different theoretical approaches, the reasoners' answers to the production task will be derived, and these predictions will be compared to the observations. The other aim is to further specify the UA, for which two experiments will be presented.

The suppression effect with the production task: Theoretical and empirical confrontation of the two approaches.

The experimental condition of BES's first experiment that is relevant to the debate provides qualitative data that consist of an analysis of the types of answers given by participants in a production task. The participants were presented with the usual three-premise set, (i. e. , for MP: *If A, C; if N, C; A*), in three different scenarios, and asked to answer the question, "What, if anything, follows?" BES' prediction is that participants should "tend to produce conclusions that refer to the additional condition that has not been affirmed or denied in the argument" (p. 354). The relevant data are the answers given by the participants who did not endorse the categorical conclusion *C* (they were slightly less than one half). There were two main categories of answer: (i) slightly more than one half of these participants answered that there was not enough information; (ii) slightly more than one quarter gave *if N, C*. The remaining categories were rare, including modal endorsements (*can-C* or *might-C*: 1%). Similar observations were made by Dieussaert, Schaeken, Schroyens, and d'Ydewalle (2000).

What interpretation did BES make of their results? They declared the *not enough information* response "consistent" with either theory (p. 357), stated that almost one third of the conclusions (the category (ii) above, plus a few *N* answers)

were as they predicted, and pointed out that few conclusions were modals, contrary to the prediction that they attributed to the uncertainty account. They concluded that "the results suggest that suppression does not arise because people doubt the truth of one of the conditionals, but instead because they doubt the validity of a simple categorical conclusion" (p. 357). But is this conclusion warranted? In order to decide, it is necessary (i) to spell out the predictions of the uncertainty account (henceforth UA), and (ii) to examine the rationale for the predictions of the counterexample account (CA).

The predictions made by the UA have been set out in Politzer and Bourmaud (2002) within the framework of mental rule theory. The general representation of a conditional *if A, C* is assumed to have the form of a *disjunctive normal form*, that is, a disjunction of conjunctive components (Politzer, 2001; 2003) similar to the form Mackie (1974) proposed for causals. Each disjunct (we need not, for the current purpose, consider more than one disjunct) is made of the conjunction of the antecedent *A* and of a number of conditions necessary for *C* to occur:

$$\{ N_1 \& \dots \& N_p \} \& A \longrightarrow C$$

The conjuncts N_1, \dots, N_p are in braces to indicate that in ordinary communication they are implicitly assumed to be satisfied: This is the first epistemic implicature mentioned earlier. We need consider only one complementary necessary condition, so that the major premise can be represented as $\{N\} \& A \longrightarrow C$. In Artificial Intelligence, some nonmonotonic logics use similar expressions for their rules of inference or for their conditional rules. For instance, in default logic (Reiter, 1980), in order to draw a conclusion *C* from a premise *P*, one spells out a list of formulae J_1, J_2, \dots called *justifications*. If *P* can be shown and not- J_1 , not- J_2, \dots cannot be shown, then *C* can be derived. The formal similarity is even closer with the formalism of circumscription (McCarty, 1980), in which an abnormality predicate $ABN(x)$ is introduced in the antecedent part of the conditional rule, which is written as: $P(x) \& \text{not-}ABN(x) \Rightarrow C(x)$. Although this kind of formalism has notorious problems with computational tractability (reviewing the knowledge base to look for possible exceptions may be intractable), it has been argued

(Politzer 2001, 2003) that human interlocutors escape this problem by virtue of the communicative principle of relevance (Sperber & Wilson, 1995). The speaker's utterance comes with a guarantee of relevance, which implies that the sentence can be exploited to produce cognitive effects, in particular to make inferences; the conditional sentence is asserted with a *ceteris paribus* assumption to the effect that the speaker does not believe or does not know that the necessary conditions such as N do not hold (this is the first epistemic implicature referred to earlier); this guarantees normality, so dispensing the hearer from the burden of verification.

We go back to the derivation of the conclusion, and we repeat here the major conditional premise: $\{N\} \& A \rightarrow C$. The additional premise *if N , C* is taken into account to the extent that the other epistemic implicature is triggered, that is, the satisfaction of N is doubted; so, N enters the context as an uncertain proposition, hence the representation of the major premise as $N \& A \rightarrow C$, where there are no more braces, for N is now explicit.

For MP, the set of premises reduces to: $N \& A \rightarrow C$; A . There are several possibilities for the reasoner.

- (i) One is to *assume* that N is satisfied, in which case $N \& A$ is true, and C follows: This accounts for the answer of one part of the participants who endorse the conclusion. One cannot exclude that some participants fail to realise that N is a necessary condition, and treat the argument without regard to the additional premise, which leads them to endorse the conclusion too. (These reasoners join those for whom the additional premise does not trigger an implicature and who, for this reason, will not treat N as doubtful and are not susceptible to the suppression effect).
- (ii) Another possibility is that participants resign themselves to the uncertain status of N and stop there for lack of information: They cannot conclude because they do not know whether or not N is satisfied and they may limit themselves to this answer. But, especially if C is in the context (as part of a question), they will achieve greater relevance by complementing this answer, stating more precisely that it is not certain whether or not C holds.
- (iii) The last possibility is to tackle the uncertainty about N : Start a suppositional line of reasoning in which N is supposed, which delivers the conclusion *if N , C* by

conditionalization. We start the derivation from the beginning, repeating the initial set of premises as 1, 2, and 3:

1. $A \rightarrow C$
2. $N \rightarrow C$
3. A

As explained earlier, premise 1 is rewritten as 4, hence the intermediate premise set 4, 5, 6 repeated here for purpose of exposition:

4. $\{N\} \& A \rightarrow C$.
5. $N \rightarrow C$
6. A

Then N enters the context (triggered by 5 as explained earlier) so that 4 is rewritten as 7, hence the new set of premises 7, 8:

7. $N \& A \rightarrow C$
8. A

The suppositional reasoning now begins:

9. N supposition
10. $N \& A$ conjunction applied to 8 and 9
11. C by MP applied to 7 and 10
12. $N \rightarrow C$ conclusion by conditionalization rule.

We now are in a position to discuss the results of BES's first experiment, and their interpretation of these results. Consider first the "not enough information" answer. For the UA, which claims that the suppression is triggered by an epistemic implicature to the effect that N might not be the case, "lack of knowledge about N " is a most directly predictable answer. However, it was not predicted by the CA. It might be explainable on the basis of post hoc hypotheses, but the authors did not attempt to do so. BES qualified the "not enough information" answer as "consistent" with both approaches, but this qualification is an overstatement with regard to the CA, which does not predict it, and an understatement with regard to the UA, which predicts it.

Consider now the next category of answers, which consisted mainly of *if N, C* conclusions. The UA predicts these logical forms with accuracy. Does the CA

explain them? The quotation of BES reproduced earlier can hardly count as a prediction, as it is almost tautological. It amounts to predicting that N (the condition "that has not been affirmed or denied in the argument") will appear in the conclusion. But since there are only three atomic propositions (A , C , and N), it is unlikely that N will not appear as one component of the conclusion. In addition, a prediction worthy of interest should contain a full logical form, not just the mention of a component such as N : Any logical form in which N appears will conform to the prediction, which in fact is hardly falsifiable. It seems that the CA has no specific basis for proposing a prediction. This is confirmed by the way the prediction is derived: It is derived from general considerations on mental models, without referring to the counterexample mechanism, suggesting that the latter actually has no role to play (unless it is assumed that it is already contained in these general principles).

In summary: The two main categories of answers observed by BES (as indicated in their Table 5), which amount to 83% of the "suppressed" answers, are predicted by the UA. In comparison, we have seen that the CA cannot make any precise prediction. Even after counting all the answers in which there is an occurrence of N (however trivial and hard to falsify this prediction may be) these answers add up to 38% of the suppressed responses. It can be concluded that, when carefully analysed, the results of BES's first experiment fail to support the CA, while they strongly support the UA.

In view of the foregoing analysis, and of the impressive support given to the UA by the literature reviewed earlier, one might wonder whether there is still a need for more data. The sceptics might base their reluctance with regard to the UA on one apparently puzzling result, namely the near absence of modal answers obtained in BES' experiment. That is, one might expect uncertain conclusions to be typically expressed explicitly by modals such as *may*, *might*, or *maybe C*, a view endorsed by BES. In spite of the fact that *we have accurately predicted* the observed pattern of answers, based on the notion that reasoners who do not assume N to be the case will answer either by expressing their lack of knowledge or (in case they enter into suppositional reasoning) by delivering a conditional response, one might still wish to have a direct demonstration of the origin of the latter category

of response, namely that hypothetical conclusions reflect participants' uncertainty. This was the aim of the first experiment.

Experiment 1

In order to demonstrate that in the standard suppression paradigm, hypotheticals are the typical way of expressing an uncertain conclusion, and that correlatively there is a limitation in the frequency of use of modals to express uncertainty, one simple manipulation suggests itself: Use a premise set in which the additional premise expresses uncertainty by way of an explicit expression of doubt, viz. , *it is not certain that N*. Since previous research (reviewed earlier) has established that in MP arguments, the uncertainty of a premise propagates to the conclusion, it will suffice to observe how people formulate their conclusion in this modified suppression task. That is, the same types of conclusions should be observed after an explicit modal additional premise as after the standard additional premise *if N, C*, if the latter does generate an epistemic implicature to the effect that it is not certain that *N* is satisfied. In brief, for the non-categorical conclusions, the modified and the standard suppression tasks should elicit similar distributions of responses.

However, one can expect one difference, which concerns the categorical *C* conclusions. In the standard task, reasoners are free to implicitly or explicitly assume the satisfaction of *N*, and consequently to draw *C* (which one third to one half of them usually do). But in the present modified task, reasoners are no more licensed to *assume N*: This is because an explicit doubt about *N* is expressed in the additional premise, which cancels the first implicature attached to the major premise; the guarantee of normality does not hold any more, so that reasoners can at most *suppose N*, and consequently few, if any, categorical *C* conclusions should be observed.

Method.

Materials and design. The experiment was presented in a paper-and-pencil format. In accordance with previous practice, the argument was preceded with a sentence stating that the premises had to be considered as true. The participants were invited to consider whether there was a conclusion that followed logically.

They were instructed (i) in case they thought that a conclusion followed, to write down this conclusion; (ii) in case they thought that no conclusion followed, to state why. There were two experimental conditions with the same single scenario which was an adaptation of the library scenario used by Byrne (1989) and subsequently by other investigators.

The *implicit condition* presented a standard suppression task with a MP as follows (translated from French):

(Major premise): *If she has a project to prepare, Nadia works in the library.*

(Additional premise): *If the library is open, Nadia works in the library.*

(Minor premise): *Nadia has a project to prepare.*

The *explicit condition* was defined by a modified additional premise: The sentence *It is not certain that the library is open* was used instead of the standard additional conditional premise.

Participants and procedure. The participants were 63 pre-University students in two suburban high schools near Paris and 47 undergraduate students of psychology untutored in logic and ignorant of the psychology of reasoning; they were randomly allocated to one of the two experimental conditions (N = 55 each). The administration took place during classes.

Results and discussion.

The responses were classified as (i) *categorical* whenever participants gave just C (unconditioned) as their conclusion (or A, or N, which happened very rarely); (ii) *hypothetical* whenever N was in the scope of *if*; (iii) *modal* whenever the conclusion (or the explanation for the estimated absence of a conclusion) contained an epistemic modality. There were very few cases of overlap between these categories; the reader is referred to Table 1 which contains the detailed surface structure of the responses within each category.

Insert Table 1 about here

We consider first participants who answered that there was a conclusion. The first main result is that in the *explicit condition*, *hypothetical* conclusions were the most frequent answers (47%); moreover, the frequency of these conclusions was about three times as high as the frequency of the *modal* conclusions (15%).

This result provides a kind of base rate for the relative frequencies of the answers when *N* is uncertain. It shows that *hypotheticals* (mostly *if N, C* and *C if N*) are given preference over *modals* to formulate the conclusion. Since the doubt expressed by the additional premise (as here in the *explicit* condition) is known to propagate to the conclusion (see the literature reviewed earlier, in particular the first experiment in Stevenson & Over, 1995), we have a direct demonstration that one of the meanings communicated by the hypothetical answer is an expression of doubt and that the modal answer (in the statistical sense) when the additional premise introduces uncertainty about *N* is conditional *on N*. In addition, the frequency of hypothetical answers in the *implicit* condition was of the same order of magnitude (51%), suggesting the essential equivalence between the additional premises used in the two conditions.

The second main result concerns participants who answered that there was no conclusion (note 4). In the *explicit* condition, only one type of answer was observed with a significant frequency (22%), namely stating the source of one's uncertainty at a metalinguistic level, such as *one does not know whether N*; the other kind of answer, the adverbial expressions such as *probably C* or *maybe C* never occurred (there was one single *maybe* occurrence but it qualified not-*N*). This shows that the rarity of this answer in the standard suppression task cannot count as evidence against the hypothesis of the processing of an uncertain premise: For this answer is as rare in the present *explicit* condition, which does contain an explicit uncertain premise to process, as it is in in the present *implicit* condition and in BES' first experiment.

These results should dispel any doubt with regard to the meaning of the hypothetical answers and any objection based on the rarity of the *may/maybe C* forms of modals. One likely reason for this rarity is conversational: Participants who are asked to *produce* a conclusion and who think that the uncertainty about *N* does not allow them to conclude about *C*, are likely to consider a modal conclusion of the type *maybe C* to be a less appropriate response *to a question test* (which they expect to be categorical) than a metalinguistic response that refers to the source of their doubt (*N*) (note 5). But of course, as mentioned earlier, participants would achieve greater relevance, although at a higher cost, in

expressing the consequence of their doubt about *N*, namely their uncertainty about *C*, which some actually did. In fact, the rarity of the *may/maybe C* conclusion seems to be entirely linked with the response format and have little to do with the inferential processes, as shown by the fact that it is the most frequent response on the standard task when this option is offered (Stevenson & Over, 2001).

Comparing the *explicit* and *implicit* conditions, two other observations are worth reporting. One concerns a difference between them: Whereas the major frequencies of response were strikingly close (*hypotheticals-yes*: 51% vs. 47%; *modals-no*: 20% vs. 22%; overall *yes*: 73% vs. 64%; overall *no*: 27% vs. 36%) there was, as predicted, a clear difference which concerns the *categorical C* conclusion: 18% of the participants gave this conclusion in the *implicit* condition but no one did in the *explicit* condition. Now, there is a correlative difference: No participant in the *implicit* condition gave a modalised *yes* conclusion but 15% of the participants in the *explicit* condition did. In brief, the response frequencies exhibited a shift from *categorical C* conclusions in the *implicit* condition to *it is not certain that C* in the *explicit* condition. This conforms to our expectations: In the suppression paradigm (here, the *implicit* condition) the implicature to the effect that the necessary condition *N* is satisfied is questioned in an indirect and subtle manner by the additional premise. This lets a proportion of participants free to maintain the implicature, and to reason in a trustful mode to produce a *categorical C* conclusion, as has been consistently observed in various investigations of the phenomenon. In contrast, in the *explicit* condition, the implicature attached to the major premise was explicitly questioned and participants were invited to reason in a distrustful mode: None of them concluded *C*, (that is, the suppression effect was total) and all those who thought that a conclusion followed expressed an uncertain conclusion either by a modal (one quarter of the time) or by a hypothetical (three quarters of the time). This observation can be taken as yet another illustration of the UA: The more explicitly the uncertainty is introduced by the additional premise, the stronger the suppression (a point already established in the literature and reminded earlier).

The other observation concerns two subcategories of hypotheticals, namely the conclusions *if not-N*, *not-C*, and *C only if N* on the one hand, and the conclusions

if N, C, and *C if N*, on the other hand. Whereas the frequency of the former subcategory was significantly greater in the *explicit* condition than it was in the *implicit* condition (13 vs. 4, chi-square = 5.63, $p < .05$), the frequency of the latter subcategory was significantly smaller (11 vs. 21, chi-square = 4.41, $p < .05$). This result is in line with our point of view: Whereas part of the *if N, C* and *C if N* conclusions may express the sufficiency of *N* for *C* and thus may constitute, if not an argument in favour of *C*, at least an agnostic standpoint with regard to *C*, the *if not-N, not-C* and *C only if N* conclusions express the necessity of *N* for *C* and thus constitute an argument against *C*. It is not surprising that such conclusions are frequent in the *explicit* condition and at the same time much more frequent than in the *implicit* condition. These fine-grained differences occurred within the main category of hypotheticals as a whole, which, to repeat, was represented about as often in both conditions, as predicted.

Finally, a few words about the *implicit* condition are in order. It must be stressed that almost all non-categorical conclusions were uncertain to various degrees and could be classified into the two predicted categories, namely modals which express lack of knowledge about *N*, and hypotheticals which express, as has just been seen, uncertainty about *N* (with the possible exception of two occurrences (4%) of *C only if N* responses, on which more will be said later). This most clearly confirms our predictions and the analysis of the results of BES' first experiment made earlier.

Taking into account these results, which add up to those reviewed earlier, we can consider the UA to be firmly established. However, some data are still missing: In the various papers reporting the suppression effect in the standard paradigm, the percentage of endorsement of the conclusion (*C*) has regularly been indicated, but that of the other two response options, viz., *not-C*, and *maybe C*, *maybe not-C*, has not. This information is important, as the UA predicts that when these options are offered, participants who do not assume *N* to be the case and who consequently should express an uncertain conclusion will choose the latter option rather than the categorical *not-C* option. This information was obtained in the second experiment, in which participants' justifications were also collected. The other main prediction was that there would be a relation between the

endorsement of the conclusion *C* and belief in the satisfaction of the complementary necessary condition *N*.

Experiment 2

Method.

Materials and design. The experiment was presented in a paper-and-pencil format. Two scenarios were used: The library scenario, as above, and an adaptation of the theater scenario used by Byrne (1989) and others, which provided the following set of premises (translated from French):

(Major premise): *If he met his girlfriend, Nicolas went to the cinema.*

(Additional premise): *If there were enough seats, Nicolas went to the cinema.*

(Minor premise): *Nicolas met his girlfriend.*

The argument was again preceded with a sentence stating that the premises had to be considered as true. This preceded the question, "What do you think of the following conclusion?" (*C* followed). Two different response formats were used, which defined two experimental conditions. Participants in the *five-option* condition ($N = 61$) were asked, "In your opinion, this conclusion is: true; slightly doubtful; doubtful; very doubtful; false?". Participants in the *two-option* condition ($N = 45$) were asked, "In your opinion, can one state that this conclusion is true?"; they had to answer by *yes* or *no*. This was a control avoiding any option formulated with epistemic terms that could have spuriously suggested the related attitudes to answer this question. Finally, all participants were asked to justify their choice in a few lines. The two scenarios were crossed with the two response formats.

Participants and procedure. The participants were 106 undergraduate psychology students who were native speakers of French, untutored in logic, and had not followed any course related to the psychology of reasoning. The administration took place in small groups. The participants had to answer only one problem; they were randomly allocated to one of the option formats as well as to one of the scenarios.

Results and discussion.

Fifty-nine percent of all the participants did not endorse the conclusion, a percentage typical of the suppression effect. The two groups behaved similarly, as 58% did not endorse the conclusion in the 2-option condition, compared to 61% in the 5-option condition. This indicates that there is no evidence that the participants were influenced by the epistemic terms contained in the response options of the 5-option condition. The breakdown for the 61% who did not endorse the conclusion in the latter condition was as follows: *False* 5%; various degrees of doubt 56% (*very doubtful* 3%; *doubtful* 25%; *slightly doubtful* 28%). In brief, participants who did not endorse the conclusion while being offered response options formulated in terms of doubt almost always used these options rather than the *false* option. The first experiment showed that when participants do not endorse the conclusion they spontaneously express their doubt (by using modals or hypotheticals). This experiment provides complementary evidence: Participants overwhelmingly preferred to declare the conclusion doubtful rather than false. Also, when the format of response offers them an opportunity to express their doubt, so indicating that this kind of answer belongs to the permitted register of answers, they make use of it.

In the first experiment, justifications were asked for only when participants did not draw a conclusion. Experiment 2 provides complementary information, as the participants had to justify their choice whether or not they endorsed the conclusion. Participants' justifications for endorsing the conclusion *C* were categorised as follows (see top part of Table 2). With respect to this categorisation, there were no differences in the distribution of the justifications between the two conditions, so that the results of the analysis were pooled together.

 Insert Table 2 about here

We consider first the 43 occurrences of endorsement of the conclusion; the percentages which follow are calculated out of these 43 occurrences. In 39% of the cases (17 observations), *C* was justified by merely repeating the MP argument or writing a paraphrase of it. Here is a typical justification: "It is said that *if A, C*; given that *A*, it follows that *C*". In 47% of the cases (20 observations), the

justification was limited to a statement of either the consequent (33%, generally preceded by "since" or "because"), or the major premise (9%), or the minor premise (5%); these short expressions may be regarded as enthymematic forms of the preceding category. The remaining 14% (6 observations) were miscellaneous justifications.

There were 63 occurrences of non-endorsement of the conclusion C ; the percentages which follow are calculated out of these 63 occurrences. The justifications were categorised in the same response categories that were used for Experiment 1. The distribution is similar, that is, a strong majority (75%) were modals (47 observations, mainly "one is not sure that N "). There were, in addition, 21% hypotheticals (sometimes worded with "on the condition that") most of which unambiguously expressed the necessity of N for C to hold. The only difference between the justifications given in Experiment 1 (for not producing a conclusion) and in Experiment 2 (for not endorsing C) comes from this subset of participants who argued in terms of the necessity of N and could be found only in the second experiment. This is understandable from an argumentation point of view: Participants who think the conclusion does not follow need to rebut it when it is presented to them. Asserting C *only if* N or *if not- N , not- C* fulfills well this purpose: This conclusion is relevant, as it enables the hearer to question the assertability of C when it is complemented with "N is unsure" in an enthymematic form. This emergence of C *only if* N and *if not- N , not C* justifications offers a likely explanation for the two occurrences of the former, and for another two occurrences of the latter (in conjunction with *if N , C*) observed in the *implicit* condition of the first experiment: Some participants who have arrived at the *if N , C* conclusion may be inclined to complement, or replace this conclusion, with the formulation of a necessary condition for purpose of justification, drawing on their metacognitive skills.

The status of the necessary condition N in participants' argumentation is crucial. The right part of Table 2 indicates whether or not N was mentioned and, in the affirmative, what its status was.

(i) all of the 30 justifications in which N was not mentioned, the 4 justifications in which N was assumed to hold (typically "because N ") and 6 out of the 7

justifications in which N was supposed to hold were cases of endorsement of C : This means that all the participants for whom the satisfaction of the necessary condition N goes without saying, and also virtually all those who assumed or supposed N to be satisfied and said so, endorsed the conclusion C .

(ii) reciprocally, all of the 47 justifications in which N was mentioned as uncertain, and 14 out of the 16 justifications in which N was mentioned as a necessary condition were cases of non-endorsement of C . In brief, the correlation between the endorsement of C (yes or no) on the one hand, and the status of N (satisfaction granted or not) on the other hand, was near to perfect (and in view of the numbers, its statistical significance need not be tested).

The fact that participants' justifications did not differ across conditions strongly supports the notion that people spontaneously approach the task in terms of uncertainty: Participants produced similar justifications whether they were asked to estimate the truth of the conclusion using the options *True* and *False* (in the two-option condition) or *True*, *False* and various degrees of doubt (in the 5-option condition) showing that they were not influenced by the format of the question.

In summary, the results are remarkably clear and in perfect conformity with the predictions of the UA. For a majority of participants, the reason for accepting C as a conclusion is based on the covert (or, less often, overt) *assumption* that the necessary condition N is satisfied, in which case the participants are dealing with an ordinary Modus Ponens. In this case, the additional premise plays no detectable role; N is implicitly joined to A and represents no obstacle to infer C once A (*qua* premise) and N (conversationally) are assumed to hold. A few participants were more explicit and took the precaution to *suppose* N , following which C was derivable in the same way as just mentioned.

The reasons for not accepting C are as clear, even though the participants' explanations were fairly concise. A majority (two-thirds) stated that one is not sure about N , about one third of which added that one is not sure about C , while a minority (about one fifth) contented themselves with pointing out the necessity of N , generally as a condition on C . In brief, virtually all these participants used an enthymematic form containing two or even only one of the three propositions

which constitute a complete justification: (i) N is necessary for C ; (ii) one is not sure about N ; (iii) therefore one is not sure about C .

Finally, the comparison of the rates of endorsement of the conclusion across experiments for participants who received the same set of premises, that is, the implicit group in Experiment 1 and all participants in Experiment 2, produces the following percentages. In Experiment 2, 41% endorsed the conclusion as true whereas in Experiment 1, 73% answered that there was a conclusion. At first sight, these two percentages differ sharply. However, the latter percentage must be broken down into unconditioned (categorical) C responses (18%) and hypotheticals (51%). While there is no doubt that the 18% who *produced* C in Experiment 1 would have endorsed C if they had participated in Experiment 2, it is clear that some of those who produced a *hypothetical* sentence in Expt 1 would have been reluctant to endorse a *categorical* C response in Experiment 2. In brief, the percentage of 41% which lies in between the strict categorical conclusion (18%) and all kinds of conclusions (73%) shows coherence between the two experiments.

General discussion

The theoretical account of the suppression effect proposed in the present study is cast in the framework of plausible reasoning: When one or more premises of a deductively valid argument are uncertain, the conclusion of the argument is uncertain. The account that is proposed is also rooted in pragmatic theory: Pragmatic analysis identifies the origin of the uncertainty as an epistemic implicature attached to the additional premise. Finally, the present account is based on a general hypothesis regarding the representation of conditionals as a disjunctive form which incorporates complementary conditions that are jointly necessary with the antecedent for the consequent to hold. These hypotheses enable one to predict, within mental rule theory, the set of possible answers to the suppression task. These predictions are clear and specific and the results of the two experiments, run with Modus Ponens arguments, have shown that they are accurate.

In Experiment 1, the standard suppression paradigm was compared with a modified one in which the additional premise *if* N , C was replaced with a modal, *it*

is not certain that N. As predicted, the same types of non-categorical conclusions were observed with similar distributions, showing that the additional premise could be substituted with an expression of doubt about *N* which produced essentially the same conclusions: This suggests the essential nature and role of the additional premise, namely introducing doubt about *N* and the major premise. In Experiment 2, the suppression paradigm was studied with an evaluation task in which participants' justifications were recorded. In line with the UA, when the conclusion *C* was not endorsed, 95% of the justifications were either modals expressing the uncertainty of the necessary condition *N* or (less often) hypotheticals expressing the necessity of *N*. When *C* was endorsed, in 93% of the cases either *N* was not mentioned or, if it was, it was explicitly assumed or supposed to hold, again in agreement with the UA.

We have applied a strict criterion to the UA to define what counts as a prediction: What should be predicted is a precise logical or linguistic form for the conclusion (e. g. , *if N, C*) or a well defined proposition that can receive paraphrastic realizations (e. g. , *one is not certain that N*). While the results based on this criterion corroborate the UA, they disprove the CA, as it does not predict the non-categorical conclusions observed in the first experiment. As noted earlier, the only prediction made by the CA is so vague that it is hardly falsifiable. It does not account for the modal conclusions (stating uncertainty or lack of knowledge) or, with any precision, for the hypotheticals without post hoc hypotheses. In addition to this empirical refutation, the CA suffers from a number of theoretical problems.

Firstly, it lacks conceptual clarity. A few times, BES claim that participants interpret the major premise *if A, C* as a *reversed conditional*, that is, as equivalent to *if C, A*. Moreover, they state that "the valid inferences are suppressed by additional requirements *because* the interpretation of the first conditional within the context of the second conditional is a "reversed conditional" rather than a conditional" (p. 364, emphasis added). This is surprising, as a reversed conditional coincides with an *only if* interpretation: It is unlikely that a premise like *If she has a project to prepare, Nadia works in the library* is interpreted by the majority of people as implying that it cannot be the case both that she does not have a project to

prepare and she studies in the library. Thus, there is a juxtaposition of the reversed conditional interpretation and of the counterexample availability (e. g., p. 350) without an explanation of how these two concepts interact to produce the suppression effect. Such a lack of integration of these two processes renders the CA obscure.

Secondly, we have seen that the CA cannot explain the hypothetical conclusions, which are pervasive. Theoretically, this is all the more problematic for the CA as this conclusion is a repetition of one of the premises, which contradicts one of the principles of mental model theory: Reasoners should not "draw a conclusion that asserts something that has just been asserted" (Johnson-Laird & Byrne, 1991, p. 22). One way for the CA to avoid this contradiction could be to consider that under the same surface structure, the meanings of *if N, C* in the premise and in the conclusion do not coincide (assuming that it could find how to explain the derivation of *if N, C*); and in order to do so, the theory should incorporate a pragmatic analysis of the premises and conclusion. From the viewpoint of the UA, the *if N, C* conclusion is derived from the supposition that *N* holds. In an argumentative background, the conclusion may, in addition, trigger an *only if* implicature of necessity (that may even be explicitly stated as was the case in Experiment 2 when participants wanted to justify that *C* is *not* true, and in Experiment 1 after participants had been presented with an explicitly uncertain *N* premise).

Thirdly, the CA does not deal with uncertainty, although expressions of uncertainty are pervasive in reasoners' answers in both experiments, as well as in the relevant literature. In order to explain the main result of Experiment 1 showing that the explicitly uncertain premise yielded globally the same kind of responses as the usual additional conditional premise, one could hypothesize that the premise *it is not certain that N* is represented by two models, *N*, and *not-N*, following which the same sets of models as in the standard task are constructed by application of the principle of pragmatic modulation. But then how to account for the effects observed in the studies referred to earlier (e. g., George, 1995; Liu et al., 1996; Politzer & Bourmaud, 2002; Stevenson & Over, 1995) in which the degree of belief in the satisfaction of *N* was manipulated. One of these effects is that the

percentage of the participants who endorse the conclusion as certain decreases with this degree of belief. Now, as the event or condition that creates the counterexample is provided in the additional premise across the various experimental conditions (that is, the various degrees of belief), it is hard to see how the decrease in the rate of endorsement of C could depend on the availability of a counterexample that is always present: This concept fails to account for the large body of results of which the suppression effect is but a particular case, showing that the CA seriously lacks generality. One possibility to remedy this would be to attribute probabilities to models; for instance, if the probability that N is satisfied is low, the model ($A \ \ N \ \ C$) would have lower probability than the model ($A \ \ \text{not-}N \ \ \text{not-}C$), and the lower N is, the less frequently people should select C . Of course, there are important points of agreement between the two approaches that have not yet been mentioned: Both approaches agree on the intervention of background conditions provided by world knowledge, and on the existence of interpretive processes that yield a logical form for the conditional premise. The debate, which focuses on how this happens, is important, as the hypothesised mechanisms lie at the heart of the psychology of plausible reasoning. The tasks that use an additional premise, like the suppression paradigm, require more elaborate treatment than the other tasks and this is why this paradigm may not be the most appropriate to study plausible reasoning and to test between opposite theories. But in spite of these difficulties, the present study has revealed the superiority of the uncertainty approach, not only by rebutting BES's claim against it, but also by showing its explanatory power with respect to the experimental results found in the literature on plausible reasoning.

In conclusion, in the suppression task, as in reasoning tasks in general, participants bring to bear all their subtle linguistic skills to deal with uncertain information and to interpret the premises: An approach to propositions in terms of belief and linguistic pragmatics are indispensable conceptual tools to provide a detailed explanation of the suppression effect.

References

- Adams, E. (1965). The logic of conditionals. *Inquiry*, 8, 166-197.
- Adams, E. (1966). Probability and the logic of conditionals. In J. Hintikka & P. Suppes (Eds.), *Aspects of inductive logic* (pp 265-316). Amsterdam: North-Holland.
- Adams, E. (1998). *A primer of probability logic*. Stanford: CSLI Publications.
- Adams, E. W. , & Levine, H. P. (1975). On the uncertainties transmitted from premises to conclusions in deductive inferences. *Synthese*, 30, 429-460.
- Bonnefon, J.-F. , & Hilton, D. J. (2002). The suppression of Modus Ponens as a case of pragmatic preconditional reasoning. *Thinking and Reasoning*, 8, 21-40.
- Boole, G. (1854/2003). *An investigation of the laws of thought, on which are founded the mathematical theories of logic and probabilities*. Amherst: Prometheus Books.
- Byrne, R. M. J. (1989). Suppressing valid inferences with conditionals. *Cognition*, 31, 61-83.
- Byrne, R. M. J. , Espino, O. , & Santamaría, C. (1999). Counterexamples and the suppression of inferences. *Journal of Memory and Language*, 40, 347-373.
- Byrnes, J. P. , & Overton, W. F. (1986). Reasoning about certainty and uncertainty in concrete, causal, and propositional contexts. *Developmental Psychology*, 22, 793-799.
- Chan, D. , & Chua, F. (1994). Suppression of valid inferences: Syntactic views, mental models, and relative salience. *Cognition*, 53, 217-238.
- Clement, C. A. & Falmagne, R. J. (1986). Logical reasoning, world knowledge, and mental imagery: Interconnections in cognitive processes. *Memory and Cognition*, 14, 299-307.
- Cummins, D. D. (1995). Naive theories and causal deduction. *Memory and Cognition*, 23, 646-658.
- Cummins, D. D. , Lubart, T. , Alksnis, O. , & Rist, R. (1991). Conditional reasoning and causation. *Memory and Cognition*, 19, 274-282.
- Dias, M. G. , & Harris, P. L. (1988). The effect of make-believe play on deductive reasoning. *British Journal of Developmental Psychology*, 6, 201-221.

- Dieussaert, K. , Schaeken, W. , Schroyens, W. , & d'Ydewalle, G. (2000). Strategies during complex deductive inferences. *Thinking and reasoning*, 6, 125-160.
- Donaldson, M. (1978). *Children's mind*. Glasgow: Fontana.
- Evans, J. St. B. T. , Handley, S. J. , & Over, D. E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 321-335.
- George, C. (1995). The endorsement of the premises: Assumption-based or belief-based reasoning. *British Journal of Psychology*, 86, 93-111.
- George, C. (1997). Reasoning from uncertain premises. *Thinking and Reasoning*, 3, 161-189.
- Johnson-Laird, P. N. , & Byrne, R. M. J. (1991). *Deduction*. Hove & London: Lawrence Erlbaum.
- Johnson-Laird, P. N. , & Byrne, R. M. J. (2002). Conditionals: a theory of meaning, pragmatics, and inference. *Psychological Review*, 109, 646-678.
- Lewis, D. (1986). Probabilities of conditionals and conditional probabilities. In D. Lewis, *Philosophical papers*. Oxford: Oxford University Press.
- Liu, I.-M. , Lo, K.-C. , & Wu, J.-T. (1996). A probabilistic interpretation of "If-Then". *Quarterly Journal of Experimental Psychology*, 49A, 828-844.
- Mackie, J. L. (1974). *The cement of the universe*. Oxford: Oxford University Press.
- McCarty, J. (1980). Circumscription – A form of nonmonotonic reasoning. *Artificial Intelligence*, 13, 27-39.
- Manktelow, K. , & Fairley, N. (2000). Superordinate principles in reasoning with causal and deontic conditionals. *Thinking and Reasoning*, 6, 41-65.
- Markovits, H. , & Vachon, R. (1989). Reasoning with contrary-to-fact propositions. *Journal of Experimental Child Psychology*, 47, 398-412.
- Neth, H. & Beller, S. (1999). How knowledge interferes with reasoning - Suppression effects by content and context. *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 468-473). Mahwah, N. J. : Lawrence Erlbaum.
- Nilsson, N. J. (1986). Probabilistic logic. *Artificial Intelligence*, 28, 71-87.
- Oaksford, M. , Chater, N. , & Larkin, J. (2000). Probabilities and polarity biases in conditional inference. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 26, 883-899.

- Oaksford, M. , & Chater, N. (2003). Conditional probability and the cognitive science of conditional reasoning. *Mind and Language*, 18, 359-379.
- Oberauer, K. , & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 680-693.
- O'Brien, D. P. , Costa, G. , & Overton, W. F. (1986). Evaluations of causal and conditional hypotheses. *Quarterly Journal of Experimental Psychology*, 38A, 493-512.
- Politzer, G. (2001). How to doubt about a conditional. In S. Benferhat & P. Besnard (Eds.), *Symbolic and quantitative approaches to reasoning with uncertainty* (pp. 659-667). Lecture Notes in Artificial Intelligence (No 2143). Berlin: Springer.
- Politzer, G. (2003). Premise interpretation in conditional reasoning. In D. Hardman & L. Macchi (Eds.), *Thinking: Psychological perspectives on reasoning, judgment, and decision making* (pp. 79-93). London: Wiley.
- Politzer, G. , & Bourmaud, G. (2002). Deductive reasoning from uncertain conditionals. *British Journal of Psychology*, 93, 345-381.
- Politzer, G. , & Braine, M. D. S. (1991). Responses to inconsistent premisses cannot count as suppression of valid inferences. *Cognition*, 38, 103-108.
- Ramsey, F. P. (1990). General propositions and causality. In *Philosophical papers*. D. H. Mellor (Ed.) (pp 145-163). Cambridge: Cambridge University Press. [originally published, 1931]
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13, 81-132.
- Rescher, N. (1976). *Plausible reasoning*. Amsterdam: Van Gorcum.
- Schubauer-Leoni, M. L., & Grossen, M. (1993). Negotiating the meaning of questions in didactic and experimental contracts. *European Journal of Psychology of Education*, 8, 451-471.
- Sperber, D. , & Wilson, D. (1995). *Relevance: Communication and cognition*. 2nd edition. London: Blackwell.
- Stevenson, R. J. , & Over, D. E. (1995). Deduction from uncertain premises. *Quarterly Journal of Experimental Psychology*, 48A, 613-643.

- Stevenson, R. J. , & Over, D. E. (2001). Reasoning from uncertain premises: Effects of expertise and conversational context. *Thinking and Reasoning*, 7, 367-390.
- Suppes, P. (1966). A Bayesian approach to the paradoxes of confirmation. In J. Hintikka & P. Suppes (Eds.), *Aspects of inductive logic* (pp. 198-207). Amsterdam: North-Holland.
- Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory and Cognition*, 22, 742-758.
- Thompson, V. A. (1995). Conditional reasoning: The necessary and sufficient conditions. *Canadian Journal of Experimental Psychology*, 49, 1-60.

Table 1. Experiment 1 (N = 110). Classification and frequency of the responses and their percentages (to be continued)

	<u>Group I</u> (N = 55) Premise set: <i>If A, C; if N, C; A</i>	<u>Group E</u> (N = 55) Premise set: <i>If A, C; it is not certain that N; A</i>
<i>Is there a conclusion ?</i>		
<u>YES: give it</u>		
Categoricals		
C unconditioned		
C	5	-
C (because N)	2	-
C (because N & A)	1	-
C (because A)	1	-
A & C	<u>1</u>	-
total	10 =	-
A	10 (18%)	(0%)
N	1 (2%)	(0%)
N	1 (2%)	(0%)
Hypotheticals		
C if N	14	7
if N, C	<u>7</u>	<u>4</u>
subtotal	21 = 21	11 = 11
if (A&N), C	3	-
A&(if N, C)	<u>-</u>	<u>2</u>
subtotal	3 = 3	2 = 2
if N, C; if not-N, not-C	2	9
if not-N, not-C	-	4
C only if N	<u>2</u>	<u>-</u>
subtotal	4 = <u>4</u>	13 = <u>13</u>
total	28 = 28 (51%)	26 = 26 (47%)
Modals		
it is not certain that C	-	7
it is not certain that C; if N, C	<u>-</u>	<u>1</u>
total	-	8 =
	(0%)	8 (15%)
Other (practical)	-	1 (2%)
Total YES	40 (73%)	35 (64%)

Table 1 (continued). Experiment 1. Classification and frequency of the responses and their percentages.

	<u>Group I</u> (N = 55)		<u>Group E</u> (N = 55)	
	Premise set:		Premise set:	
	<i>If A, C; if N, C; A</i>		<i>If A, C; it is not certain that N; A</i>	
<i>Is there a conclusion ?</i>				
<u>NO:</u> <i>Why ? Because:</i>				
Modals				
<i>don't know whether N</i>	7		7	
<i>don't know whether N</i>				
<i>∴ don't know whether C</i>	3		1	
<i>maybe not-N</i>	1		-	
<i>it is not certain that N</i>	-		2	
<i>Pr(not-N) = 1/2</i>				
<i>∴ one cannot know</i>	-		1	
<i>we talk in the conditional</i>	-		<u>1</u>	
total	11	=	12	=
			11	(20%)
			12	(22%)
Miscellaneous				
<i>not enough information</i>	1		-	
<i>if (A & not-N), not-C</i>	1		-	
<i>neg (not-N because A)</i>	-		1	
<i>can't deduce anything</i>	-		1	
<i>premises are incompatible</i>	-		<u>1</u>	
total	2	=	3	=
			2	(3.5%)
			3	(5%)
Other (practical)			2	(3.5%)
			5	(9%)
Total NO			15	(27%)
			20	(36%)

Table 2. Experiment 2 (N = 106). Classification and frequency of the justifications.
(to be continued)

Was C endorsed ?		Mention of N					NO
		YES, as					
<u>YES : 43</u>		uncer- tain	neces- sary	a supp- osition	an assu- mption	other	
<u>justification</u>							
17	Modus Ponens	15	unqualified				15
		1	assuming N				
		1	supposing N				
4	Major premise	3	unqualified				3
		1	assuming N				
2	Minor premise	2					2
14	Consequent C	2	<i>C because N and A</i>				2
		2	<i>C because A</i>				2
		1	<i>C since M if N</i>				1
		3	<i>C if N</i>				3
		1	<i>if A or if N, C</i>				1
		5	<i>since A, C</i>				5
3	<i>C is consistent with premises</i>						3
1	<i>A or N is sufficient</i>						1
1	<i>N is necessary</i>						1
<u>1</u>	<i>C if and only if N</i>						<u>1</u>
Total	43	0	2	6	4	1	30

Table 2 (continued). Experiment 2. Classification and frequency of the justifications.

Was C endorsed ?

NO : 63

Justification

47 **Modals**

- 42 *one is not sure that N* unqualified 23
 - after major premise or MP 6
 - \therefore *not sure that C* 13
- 4 *it is possible that not-N* 4
- 1 *unexpected event preventing N* 1

13 **Hypotheticals**

- 5 *C only if N* 5
- 3 *if not-N, not-C* 3
- 1 *C except if not-N* 1
- 2 *C on the condition that N* 2
- 1 *MP on the condition that N* 1
- 1 *C if N* 1

3 **Other**

- 2 *N is necessary* 2
- 1 *N is insufficient* 1

Total	63	47	14	1	0	1

Mention of N					
YES, as					NO
uncer- tain	neces- sary	a supp- osition	an assu- mption	other	

Footnotes

1. Another explanation was also envisaged, based on the hypothesis of the conversion of the additional premise. With hindsight, it does not seem to be correct: There are experimental data against it (Bonnefon & Hilton, 2002; Byrne, Espino & Santamaría, 1999, Experiment 2).
2. The demonstration is trivial in both cases. For conditional probability, we have $p(C/A) = p(A\&C) / p(A) = p(C)$ since $p(A) = 1$. For the material conditional: $p(\text{if } A, C) = p(\text{not-}A \text{ or } C) = p(C)$ since $p(\text{not-}A) = 0$.
3. BES (pp. 349 to 353) exemplified their theoretical claim using MP only. For this reason, and also for the sake of simplicity, only MP arguments will be experimentally considered in the present study.
4. A category of answer (called "practical" in Table 1) appeared with a frequency of 7%. It was given by the High School students (and this was the only difference between these and the University students). It consists of higher order conclusions based on practical considerations, such as "she will go and check whether the library is open" or "the library is not the only place to work".
5. In an experimental setting, and especially when participants are students, the choice of the answers is partly guided by tacit conventions about what is acceptable in terms of informational content, as emphasized by developmental psychologists (Donaldson, 1978) and by educationalists who have applied research initially conducted in the classroom to the laboratory (Schubauer-Leoni & Grossen, 1993).