Journal of Psycholinguistic Research, Vol. 20, No. 2, 1991

Are Conjunction Rule Violations the Result of Conversational Rule Violations?

Guy Politzer[1] and Ira A. Noveck[2,3]

Accepted November 1, 1990

This study is concerned with the conjunction error as reported by Kahneman and Tversky (1982). In the prototypical problem that manifests this error, the Linda *problem, subjects are presented with a description of a 31-year-old liberal progressive named Linda and asked to rank in order a number of events concerning Linda today in terms of their likelihood. The three events of interest are* Linda is a bank teller *(event* A*),* Linda is active in the feminist movement *(event* B*), and* Linda is a bank teller and active in the feminist movement *(event* A and B*). Subjects typically ranked* Linda is a bank teller and active in the feminist movement *(*A and B*) higher than* Linda is a bank teller *(*A*). This judgment violates a fundamental law of probability, namely, that* p *(*A and B*)* $\leq$ p*(*A*). Kahneman and Tversky argue that this response tendency is a result of the* representativeness *heuristic which causes subjects to ignore logical considerations of set relations and to rank the choices in terms of their resemblance to Linda's description. On the other hand, we argue that the task demands compel subjects to interpret the choice* A *as* A and not-B*. When subjects make the unusual comparison between* A *and its subclass* A and B *in the context of a description of Linda that is largely irrelevant to the task's solution,* Linda is a bank teller *becomes* Linda is a bank teller and is not active in the feminist movement*. With this interpretation, subjects' typical response patterns no longer can be considered in contradiction to logical principles. One experiment is presented that uses three French language* Linda*-type problems in two conditions: one in which the expected implicature is blocked and one in which the implicature is encouraged to reappear. These are both compared to a French version of the* Linda *problem. The structure of the* Linda *problem is discussed in light of conversational principles, particularly relevance (Sperber and Wilson, 1986) and the maxim of quantity (Grice, 1975).*

[1] CNRS and Université de Paris VIII.

[2] New York University.

[3] Requests for reprints should be addressed to Ira Noveck, Department of Psychology, New York University, 6 Washington Place, 8th floor, New York, New York 10003. The authors express appreciation to Martin Braine, George Davidson, Brooke Lea, David O'Brien, Mark Samuels, Dan Sperber, and the members of CREA for their thoughtful comments and suggestions.

0090-6905/91/0300-0083$06.50/0 © 1991 Plenum Publishing Corporation

One of the most striking effects reported in the probability judgment literature is the conjunction error (Kahneman and Tversky, 1982; Tversky and Kahneman, 1982; Tversky and Kahneman, 1983) that refers to the violation of a fundamental law of probability, $p(A \text{ and } B) \leq p(A)$. Given two events, A and B, the probability of the event A *and* B is necessarily less than or equal to the probability of either of its constituents alone. These authors have demonstrated that, in a particular type of problem that requires subjects to compare several events in respect to their likelihood, the majority of subjects attributed higher probability ratings to the item depicting event A *and* B than to the item depicting event A alone. For instance, in what has become a kind of prototype, the *Linda* problem, subjects are presented with a brief description (hereafter referred to as the problem's *stem*) of a woman's character that portrays her as concerned with issues of discrimination, social justice, etc., and the subjects' task is to rank order a number of assertions, three of which are the following:

1. Linda is a bank teller (A).
2. Linda is active in the feminist movement (B).
3. Linda is a bank teller and is active in the feminist movement (A *and* B).

The percentage of subjects who rank the probability of Linda being a bank teller and active in the feminist movement (event A *and* B) higher than that of Linda being a bank teller (event A) is typically around 85%.

Tversky and Kahneman (1982, 1983) have proposed an explanation of this effect based on the hypothesis that people do not solve the task by applying extensionsal reasoning, but rather by using the intuitive heuristic they call *representativeness*. They define representativeness as an assessment of the degree of correspondence between a model and an outcome. In the Linda problem, being an active feminist is representative of the character description, whereas being a bank teller is nonrepresentative. They argue that the conjunction of a representative feature (feminism) and a nonrepresentative one (Linda's being a bank teller) as presented in item 3 is judged to be more likely than item 1, which contains only the nonrepresentative component of item 3, because item 3 has at least one of the model's representative features whereas item 1 has none. Tversky and Kahneman (1983) showed that the conjunction rule violation was committed in experiments where subjects were asked to make judgments in domains as varied as social activity, symptoms of disease, sports

outcomes, and dice throw outcomes. Further work has confirmed the generality of the effect (Leddo, Abelson, & Gross, 1984; Locksley & Stangor, 1984; Wells, 1985).

There are a number of problems with the explanation of the effect as presented by Kahneman and Tversky (1982) and Tversky and Kahneman (1983), and there are reasons to believe that the conjunction effect might be partly artifactual. We argue that the task achieves its effect because the pragmatic demands of the task encourage subjects to reinterpret the presented materials before carrying out the task's required extensional analysis. A linguistic explanation of subjects' performance on the Linda problem is supported by Fiedler's (1988) finding that the error rate decreases dramatically when subjects are asked to rank in order the listed events with respect to their *frequency* as opposed to their *probability*. This shows that the effect is partly linked to the interpretation of the word *probability* and demonstrates that there exists at least one subtle linguistic factor that discourages violations of the conjunction rule. In the present work, we examine the pragmatics of the overall paradigm, focusing on how the joint presentation of items 1 and 3 leads subjects to mistakenly represent item 1 as event *A and not-B*, that is, as:

4. Linda is a bank teller and is not active in the feminist movement (*A and not-B*).

With this interpretation of item 1, subjects' typical response patterns are not in contradiction to logical principles. Below, we shall review the task's format and later specify the conditions under which subjects can be led to misconstrue item 1 as 4.

PRAGMATIC ANALYSIS OF THE LINDA PROBLEM

The most conspicuous anomaly of the Linda problem lies in its question format, which contains two items that do not have the same degree of specificity: One item contains one event, event *A*, as represented in item 1, and one contains an event that conjoins *A* and *B*, as represented in item 3. Asking subjects to compare nested classes (where event *A and B* is included in event *A*) actually has no ecological validity. This makes the Linda problem pragmatically strange.

In everyday life, the felicity conditions of a question are typically respected. These include the conditions under which one does not expect to receive a question whose answer one assumes the questioner can work out. For instance, for one to ask "How many children do you have?"

after having heard a parent say "I have one boy and one girl" is infelicitous because the answer to such a question is inferable by a mathematical operation that the parent would assume the questioner is capable of carrying out. In the same respect, the Linda problem does not respect felicity conditions. The Linda problem asks the subject to make an inclusion judgment in which a class and one of its subclasses are compared. Since the question seems trivial, subjects are forced to reinterpret it. Linguistic intuition suggests that questions that compare a superordinate class A and a subordinate class A_1 seem odd unless the hearer can interpret A as referring to A *and not-*A_1. This intuition is substantiated theoretically and empirically.

Theoretically, three explanations are available that account for the anomaly in comparing a class A to its subclass A_1 and for the mechanism that would reinterpret the comparison in order to make it more sensible. These will be only outlined below since it is not the aim of this paper to test between them. One approach is Gricean (Grice, 1975) and hinges on the quantity maxim. It requires that, just as one should not give too little or too much information, one should not ask too little or too much. Upon being asked a question whose answer the questionee estimates to be trivial, the questionee reasons that the question presented cannot be the intended one and makes an adjustment that transforms the comparison. Thus an implicature is generated in order to lend sensibility to the comparison, that is, A *vs.* A_1 becomes A *and not-*A_1 *vs.* A_1. In terms of the Linda problem, item 1 vs. item 3 becomes item 4 vs. item 3. A second approach can be described in terms of the *rationality* principle (Kasher, 1976). This principle pertains to the notion that each person assumes her interlocutor to be rational. According to this view, the questionee will assume that the questioner could not intend such a trivial comparison and the questionee is compelled to search for another referent for the item depicting event A. A third account could be formulated in terms of relevance (Sperber and Wilson, 1986) that values an utterance in terms of its ability to work out contextual implications. In problems such as the Linda problem, the question posed is irrelevant because the answer already has been stipulated contextually. The task misleads subjects into supposing that the context, as provided by the stem, is relevant when, in fact, it is actually irrelevant to arriving at a correct ranking. The result is that the questionee searches for new interpretations of the question. In the case of the Linda problem, the questionee searches for a new referent for event A. All three of these proposed discourse principles concur with the notion that, when an inclusion question is asked in normal conversation (that is, a question that requires a comparison

between a superdinate A and subordinate A_1), the questionee usually performs a normalization by which event A receives a new referent (namely, the class complement of A_1, *A and not-*A_1). When A_1 is defined as being different than A by a qualification (where A_1 = *A and B* and B is a qualifier of A), the pragmatic normalization amounts to replacing A with *A and not-B*.

Normalizations of this sort have been described in order to explain Inhelder and Piaget's (1964) well-known class inclusion problem that also involves comparisons between a class and one of its subclasses (Politzer, 1989). In Inhleder and Piaget's study of the development of concrete operatory thought, they asked their young subjects questions of the type "Are there more daisies or more flowers?" while the children were viewing, for example, seven flowers, four of which were daisies. Typically, the children said that there were more daisies. In this situation, it has been shown that a number of disambiguating linguistic manipulations result in increased correct performance by clearly making *flowers* refer to *all the flowers*. The disambiguation encourages the children to make the comparison, albeit a trivial one, between the subclass *daisies* and the class *flowers* instead of transforming the comparison into that of two subclasses, *daisies* and *flowers-that-are-not-daisies* (see, e.g., McGarrigle, Grieve, & Hughes, 1978; Shipley, 1979).

Several investigators have tested the hypothesis that a similar normalization occurs in the Linda problem. In their investigation of the conjunction effect, Yates and Carlson (1986) found direct evidence for cases of *A and not-B* interpretations in their subjects' protocols, as did Agnoli and Krantz (1989) in one of their problems. Tversky and Kahneman (1983) themselves considered the possibility that *Linda is a bank teller* might be interpreted as *Linda is a bank teller and is not a feminist*. Again, note that if such an interpretation takes place, subjects' responses no longer can be considered fallacious. Below we review experiments that investigate whether subjects make such an interpretation.

Tversky and Kahneman (1983) used three manipulations to test the hypothesis that subjects interpret item 1 as 4. In the first one, they presented subjects with the two items, 1 and 3, and rather than ask subjects to indicate which of the two was more probable, asked them to rate the probability of each item. They still found that 82% of subjects assigned a higher rating to the item represented by event *A and B* than to the item representing event A, a percentage equivalent to that obtained by the item-ranking technique. This result is not surprising. If an *and not-B* implicature is indeed tacked on to the item represented by event A in the standard task, then it would be expected in this revised task as

well. According to a conversationalist perspective, there is no reason to expect a probability *rating* task to give a result different from a probability *ranking* task. Tversky and Kahneman's second manipulation seems more appropriate at first sight. They replaced the *A* item with *A whether or not B (Linda is a bank teller whether or not she is a feminist)* in order to "emphasize the inclusion of *A and B* in *A*" (p. 299). Unfortunately, this phrase also introduces uncertainty among subjects as to whether Linda is indeed a feminist. Actually, the authors found a high rate of conjunction errors (57%) but did not make any comparisons to the standard task. If we use the value of 82% as a measure for the rate of errors in the two-item form of the problem in their first manipulation, one finds that there is a 25% decrease in the error rate. This reduced error rate demonstrates that for a sizable minority of subjects, the phrase *whether or not* may very well have canceled the implicature *and not-B*. (In a similar manipulation, Locksley and Stangor, 1984, failed to observe any effect in which they used two conjunctions in the same sentence. Their wording, however, became very awkward when *whether or not* was introduced). Finally, in a third experiment, a between-subjects comparison was performed in which different subjects rated the probability of *A*, *A whether or not B*, or *A and B* (in each case against common filler items). The results showed that the three items were rated in an increasing order of probability. This finding will be discussed in the conclusion, but note that such a manipulation does not address the normalization process that the conversationalist hypothesis claims to occur in the context of items 1, 2, and 3 appearing together, as they do in the standard task.

Another attempt at checking for a possible misinterpretation of the item depicting event *A* was made by Morier and Borgida (1984). They compared two forms of the problem, the standard form and what they referred to as a *debiased* form. In the debiased form, the additional item *A and not-B* was presented in order to make explicit the *and not-B* implicature and, by contrast, hopefully render the item depicting event *A* alone noncommital with respect to event *B*. For the Linda problem, the subjects received either a standard version with seven items or a debiased version that contained the additional item, *Linda is a bank teller who is not a feminist.* Subjects also received another problem that described a man, Bill, who likes to spend time with his wife. The standard version of the *Bill* task required subjects to estimate the probability of:

5. Finding Bill and Bill's wife at home.
6. Finding Bill at home.
7. Finding Bill's wife at home.
8. Finding at least Bill at home.

The debiased version of the Bill problem made use of the same four items, 5–8, plus two others, viz., finding Bill at home alone and finding Bill's wife at home alone. Whereas there was a significant decrease in the fallacy rate between the standard and debiased versions of the Bill problem, from 77% to 49%, no improvement was obtained for the Linda problem. The results of this manipulation are inconclusive.

Finally, Agnoli and Krantz (1989) carried out a manipulation specially designed to test the conversationalist hypothesis, that is, to test whether item 1 is interpreted as 4 when it is presented along with item 3. They carried out a between-subjects study that compared two slightly different sets of items in the Linda problem. One set of items contained items 1 and 3, plus some filler items, and the other modified set contained items 4 and 3, together with the same filler items. It was the experimenters' intention to force the conversationalist interpretation in the former treatment. As a dependent measure, they used the difference between the subjects' rankings for item 1, event *A*, and item 3, event *A and B*, in the former condition and the difference between the subjects' rankings for item 4, event *A and not-B*, and item 3, event *A and B*, in the second condition. In their experiments, a positive rank difference represents a logically correct response because item 1 is ranked higher than item 3, and a negative difference represents an error of conjunction because that means item 3 was ranked higher than item 1. According to Agnoli and Krantz (1989), if subjects interpret item 1 as item 4 then the rank difference of interest for each of the two groups should be comparable. Below we review their findings and show that an assumption underlying this prediction is questionable.

First of all, they found that the modified set gave rise to fewer positive rank differences in comparison to the more standard set, as both the conversationalist hypothesis and the heuristicist approach predict. Second, they observed that the modified set gave rise to more negative rank differences; they argued that this is predicted only by the heuristicist approach. The argument is based on the assumption that item 4, when presented as an item, is viewed by subjects as being even more dissimilar to item 3 than is item 1 to item 3. As we stated above, they argued that the conversationalist point of view should predict that the comparison between item 3 and item 4, the interpretation that subjects have for item 1 anyway, should yield the same results as a comparison between item 3 and item 1. However, this prediction does not follow from the heuristicist approach without an additional assumption left implicit by the authors. The assumption is that there exists a common behavior for the entire population such that all subjects fall prey to the *and not-B* impli-

cature. This is not a warranted assumption since all the experimental data indicate that there is a proportion of subjects, ranging from 10% to 25%, who respond correctly to the standard items. When presented with the modified set of items, the one containing items 3 and 4, those subjects who would have solved the more standard problem using the appropriate extensional analysis were prevented from doing so. Thus, for the modified problem, those reasoners who would have solved the original problem anyway were forced to join ranks with those subjects who normally would have given the incorrect, modal response. This analysis is supported by the fact that the increase in negative rank differences, not reported by the authors but derivable from the diagrams, seems to be only around 10%. We conclude that Agnoli and Krantz did not disprove the conversationalist hypothesis.

BLOCKING AND MANIFESTING THE IMPLICATURE

In view of the difficulty in finding an appropriate alteration of the items that would cancel the implicature within the original version of the task or that would render the implicature explicit, we felt that a semantic technique that blocked the implicature would be equally effective in demonstrating that subjects interpret item 1 as 4 and thereby are not violating the conjunction rule. This aim could be met by using items whose semantic relation is such that no implicature is generated. This is the case with some action concepts (as opposed to numerable classes).

Consider an event E that is logically entailed by two outcomes, events O_1 and O_2, as illustrated by items 8, 9, and 10 below:

8. Roland took an exam (E).
9. Roland failed an exam (O_1).
10. Roland passed an exam (O_2).

The consideration of nonconjunctive nested items as experimental materials is legitimate in the sense that Tversky and Kahneman's (1983) paper addressed the problem of the extension rule, viz., the *inclusion* rule: If the extension of A includes the extension of B, then $p(A) \geq p(B)$. The conjunction rule as it is used in the Linda problem is but a particular case of this. In the Linda problem, event A *and* B is necessarily included in event B, and the same principle can be expressed by the conjunction rule $p(A \text{ and } B) \leq p(B)$. Consequently, the use of items that have a relation of inclusion between them is appropriate in testing for the ex-

istence of a general inclusion rule violation in subjects' extensional reasoning.

It is hypothesized that no implicature is generated for item 8 for the following reason. The implicature is triggered when one item is compared to a second that is identical except for the addition of a qualifier. In the Linda problem, two items depicting event *A* and event *A and B* are compared and the explicit absence of the qualifier *B* in the first item amounts to the implicature of the qualifier *and not-B* in the item depicting event *A* alone. No qualifier is apparent in items 9 and 10 because in each of these only the qualifying phrase is presented and the qualified phrase in each is presupposed. In the present study, the kind of material exemplified by items 8, 9, and 10 above was used in one condition to test the hypothesis that when the implicature is blocked in this manner, the rate of conjunction rule violations will decrease or disappear. But this manipulation would not suffice to demonstrate the effect of blocking the implicature because it also increases the salience of the inclusion relationship.

Clearly the inclusion relationships among items 8, 9, and 10 are more salient than those among items 1, 2, and 3 on the standard problem. And it is assumed that, *ceteris paribus*, the more salient the inclusion relationships among the items, the less frequent will be subjects' errors when ranking the items. In order to demonstrate that the implicature can be generated even in a setting whose inclusion relationships are relatively salient, a second condition was introduced.

It was hypothesized that if the sheer comparison of the two events *A* and *A and B* in the Linda problem could generate the implicature *and not B* for event *A*, then explicitly stating event *E* plus the word *and* for the semantically entailed event in the outcomes O_1 and O_2 would do the same and tend to reintroduce the error. Thus, items 8 through 10 above would become:

11. Roland took an exam (*E*).
12. Roland took an exam and failed (*E and* O_1).
13. Roland took an exam and passed (*E and* O_2).

This manipulation was designed to demonstrate that, despite the cues that make the set/subset relationships more salient, the explicit repetitive mention of event *E* plus the word *and* in items 12 and 13 would reintroduce the pragmatic peculiarity that is found in the Linda problem. Compared to the problem in the first condition, it was expected that in a problem using items of the sort found in 11 through 13 subjects would be more susceptible to interpreting event *E* as event *E and not* O_1 or *E*

and not O_2 (depending on the content in the stem) and consequently should make more inclusion rule violations.

Predictions about subjects' performance on these versions of Linda-type tasks derived from the heuristicist approach are the following. Since all the scenarios used by Tversky and Kahneman (1983) in their study of the conjunction effect gave rise to a rate of rule violations ranging from 70% to 100%, such high rates should be expected for all inclusion-type problems, regardless of whether the presupposition and the word *and* are explicitly stated. According to Tversky and Kahneman, subjects should select the item that has the highest degree of representativeness. That is, assuming the items presented above concerning Roland were presented in the form of a Linda-type reasoning problem and that the presented stem informs subjects that Roland is talented and determined, subjects should rank item 10 highest in the first condition and rank item 13 highest in the second, regardless of the manner in which the items are presented. In both conditions, the degree of representativeness is the same. In contrast, the conversationalist analysis developed above leads to the following prediction: With respect to the rate of conjunction rule violations on the standard Linda problem, there should be a lower rate on both types of inclusion problems. This prediction stems from the hypothesis that subjects are able to appreciate the implications of the relationship between a set and its subset and more able to do so when this relationship is salient. Furthermore, there should be a lower rate of errors for the problems that do not contain the explicitly stated presupposition and the word *and* among the items than for the problems that do. This prediction is derived from our main hypothesis that subjects' poor performance on Linda-type problems is, at least partly, due to the conversational implicatures subjects make as a result of the manner in which items are presented and that in the absence of an implicature the rate of errors should decrease. In summary, in terms of rates of error the following order is predicted: standard problem > inclusion problems with the explicitly stated presupposition plus the word *and* > inclusion problems without the explicitly stated presupposition or the word *and.*

METHOD

Subjects. One hundred-and-thirty-two first year students of computer science from an *Institut Universitaire de Technologie* near Paris served as subjects. All of them had taken a course in elementary probability.

Materials. Seven problems were prepared, of which each subject

received four. Of these four, one was the French language equivalent of Kahneman and Tversky's (1982) Linda problem, referred to as the *Helene* problem, having the standard, relevant three items:

14. Helene is a bank teller.
15. Helene is active in the feminist movement.
16. Helene is a bank teller and is active in the feminist movement.

The Helene problem, which can be found in the appendix, was used both in order to calibrate the population and to serve as a control. As a calibration problem it was used to verify that it would produce the usual high rate of error among French speakers. As a control, this problem was used to verify the main hypothesis, namely, that it would give rise to a higher rate of error than the other problems. The other presented problems were based on three character descriptions novel to this paradigm and were prepared in one of two formats. These novel character descriptions and the two versions (i.e., item formats) in which they were prepared will be discussed in greater detail below. In every problem, the subjects' task was to rank in order three presented items concerning events or predictions from 1 (for *most likely*) to 3 (*least likely*).

Novel Character Descriptions. Six versions of Linda-type problems translated into English can be found in the appendix. They were prepared based on descriptions of three persons and in two item formats. They contain three items that were designed such that the event depicted in one item served as a presupposition for the other two. The *Daniel-implicit* problem can serve as an example. Subjects were told that Daniel was a bright high school student and were given the following three predictions concerning his performance in his pursuit of further study:

17. Daniel entered Medical School.
18. Daniel dropped out of Medical School for lack of interest.
19. Daniel graduated from Medical School.

As explained above, each problem was also prepared in an explicit format, i.e., the presupposition (that Daniel entered Medical School) for items 18 and 19 was made explicit along with the word *and*. To continue with the Daniel problem as an example, the Daniel-explicit problem's items became:

20. Daniel entered Medical School.
21. Daniel entered Medical School and dropped out for lack of interest.
22. Daniel entered Medical School and graduated.

The other two novel character descriptions were also prepared in

implicit and explicit formats. One, the *Lendl* problem, was inspired by tennis scenarios from Kahneman and Tversky's (1983) *Borg* problem and Yates and Carlson's (1986) *McEnroe* problem. The *Edouard* problem differs slightly from the other novel character descriptions in that the items presented in it reflect varying degrees of applied violence. Nevertheless, the predictions concerning Edouard present a representative item (*he will not hesitate to kill the aggressor*) that has a semantic constituent which serves as another item (*he will not hesitate to use violence against the aggressor*). Hereafter, the problems will be referred to as *Helene, Lendl-implicit, Lendl-explicit, Daniel-implicit, Daniel-explicit, Edouard-implicit,* and *Edouard-explicit.*

Procedure. Subjects were run in one group. All subjects received booklets containing an introductory page, the Helene problem, plus one of the two versions concerning each of the three novel character descriptions. The instruction page said, "You will find several descriptions concerning real or imaginary persons," and further said that there would be a question to answer concerning each of these persons.

Design. There were 12 orders as a result of the following criteria: (1) The Helene problem appeared either first or last to test for possible effects of training; (2) among the three novel character descriptions presented, (a) order was counterbalanced across character descriptions and (b) the versions across the three problems were alternated such that the first and third character descriptions used the same item format, leaving the item format of the second character description in each booklet unique. See Table I for the actual design of the experiment.

RESULTS AND DISCUSSION OF RESULTS

A subject's response was considered to be in error when either of the two items that represent the nested events, e.g., item 18 or 19 in

Table I. Experimental Design: Presentation Order of the Problems[a]

I	Lendl-explicit	Edouard-implicit	Daniel-explicit
II	Lendl-implicit	Daniel-explicit	Edouard-implicit
III	Edouard-implicit	Lendl-explicit	Daniel-implicit
IV	Edouard-explicit	Daniel-implicit	Lendl-explicit
V	Daniel-explicit	Lendl-implicit	Edouard-explicit
VI	Daniel-implicit	Edouard-explicit	Lendl-implicit

[a]*Note*: Each line presents only the order of the problems that used novel character descriptions. Besides the three problems listed in each order, all subjects received the Helene problem, the French equivalent to the Linda problem. For each of these orders, half the subjects received the Helene problem first and half the subjects received it last.

Daniel-implicit and item 21 or 22 in Daniel-explicit, was judged to be more likely to occur than the presupposed event presented alone, e.g., item 17 in Daniel-implicit and item 20 in Daniel-explicit. Table II presents the percentage of inclusion rule violations grouped as Helene, -explicit, and -implicit calculated over the 132 subjects and all the problems. No significant effect for training was found. Accordingly, the results were collapsed. The percentages of conjunction rule violations are in the order predicted, that is: Helene > -explicit > -implicit.

In order to test for statistical significance, the Jonckheere trend test for a two-category response variable corrected for continuity (Leach, 1979) was applied between subjects. The three groups were constituted by selecting the subjects who received the Helene problem first, an -implicit problem first, and an -explicit problem first, thus making use of all 132 subjects in a single measure. The trend was highly significant ($Z = 3.43, p < .0002$).

In order to make finer tests for significance between the three conditions, the results from those problems using novel character descriptions that had a unique format in each booklet were extracted and used as between-subject data (see Table I). That is, the responses to the -implicit version problems gathered from the six booklets containing a single -implicit version were compared to the data from the Helene problems that were in the booklets containing one -explicit version. Likewise, the responses to the -explicit version problems gathered from the six booklets containing a single -explicit version were compared to those from the Helene problems that were in the booklets containing a single -implicit version. Finally, the two sets of data coming from booklets containing a single -implicit and -explicit version were compared to one another. The results of the trend test were confirmed by chi-square comparisons performed on groups selected in the manner indicated above. As predicted, the percentage of inclusion rule violations was significantly lower among the -implicit problems than the Helene problem [$\chi^2(1) = 44.1, p < .0005$, one tailed]. The percentage of inclusion rule violations

Table II. For Each Problem, the Percentage of Subjects Who Ranked the Options in a Manner That Reflected a Conjunction Rule Violation

	Problems			Total
Helene				77.3
Novel characters	Daniel	Edouard	Lendl	
-Implicit	31.3	33.8	21.5	28.9
-Explicit	53.1	53.0	16.6	40.8

was significantly lower among the -explicit problems than among the Helene problems [$\chi^2(1) = 11.1$, $p < .0005$, one-tailed]. Finally, as expected, there were fewer errors in response to the -implicit problems than there were to the -explicit problems [$\chi^2(1) = 4.07$, $p < .025$, one-tailed].

The foregoing analyses in which the rates have been collapsed across scenarios slightly misrepresent the results due to the Lendl task, a task whose scenario did not produce the same effects as the other two novel scenarios. The rate of errors on both Lendl problems (around 20%) approaches floor; the logical relationship among its items seems to have been so compelling that the explicit wording plus the word *and* manipulation was ineffective. The rate of error for the Daniel-implicit and Edouard-implicit problems was around 30%, showing that only a minority of subjects did not respect set relations in solving these problems. It is with these two scenarios that the introduction of the explicitly stated presupposition along with the word *and* was effective as predicted.

GENERAL DISCUSSION

This experiment clearly bears out pragmatic considerations that prior analyses of the standard task have not thoroughly investigated. Though the error rate was not completely eliminated, the reduced error rate among all the novel character descriptions demonstrates that the majority of the subjects were indeed sensitive to the logical structure of those problems' included classes. Moreover, the mere presence of those items whose semantic constituents were explicitly stated along with the word *and* increased the error rate, thus demonstrating that subjects were sensitive to the pragmatically peculiar demands attached to directly comparing an event E to an event E and its possible outcomes.

The heuristicist viewpoint identifies representativeness as a sufficient condition for the modal response on the Linda problem. However, the results from the experiment carried out above bring this claim into question. The results of the first part of the manipulation, where nested items were used and more than 70% of the subjects answered correctly, cast doubt on Tversky and Kahneman's (1983) claim that "people are not attuned to the detection of nesting among events, even when these relations are clearly displayed" (p. 304). Likewise, the result of the second part of the manipulation, where the explicit wording plus the word *and* among the items was shown to be responsible for a 20% increase in the rate of error, questions the same authors' view on the

proper role of *and* ("the conjunction fallacy is not restricted to esoteric interpretations of the connective *and*," p. 303). It can be argued that they found evidence of the error in their *die-throw* situation that did not involve explicit *and* wording, but this task differs from the others in that the inclusion relation is not defined on a semantic basis and may therefore remain unnoticed. The solution to the die problem requires some knowledge of probability theory (such as calculating the probability of a sequence of events). In addition, the die situation violates another pragmatic convention that may be crucial to the task and to which we now turn.

There is a convention in the educational system that a test's *givens*, though not necessarily thoroughly informative, contain some measure of relevance to the required judgment. In other words, the information in a stem is typically relevant to the question asked. The problems used in the standard Linda-type studies, however, contain a scenario that is not necessary to solve the problem. All the information relevant to the problem's solution is contained in the relationships among the items (i.e., event *A and B* is necessarily included in event *A* regardless of the story in the stem). In brief, the Linda problem's standard format violates the conventions between tester and testee and, as such, the task can be described as deceptive. Consequently, we are not in agreement with Tversky and Kahneman who, while considering the Gricean (1975) principles, assert that the conversational analysis simply does not apply to numerical estimates and choices of bets (Tversky & Kahneman, 1983, p. 312). They overlooked the violation of tester–testee conventions in their problems, including the die throw problems. In these, the information provided, in terms of the ratio of greens and reds, is not relevant to the correct solution.

All the problems used to study the conjunction effect are designed with (i) a general frame that sets up a topic in which the items will be meaningful, e.g., introducing Linda as an adult, (ii) irrelevant pieces of information in the sense that they are not necessary to find out the extensionally correct solution, e.g., Linda's character, opinions, etc., and (iii) item comparisons that are pragmatically unsettling. Due to the tester–testee convention discussed above, the subject assumes that the pieces of information described in (ii) are indeed necessary for the task's solution. Consequently, the subject's strategy will be to work out interpretations of the items that can render them relevant to the stem of the problem. The subject can succeed in doing this by interpreting the item depicting event *A* as event *A and not-B*. In contrast, the subject who solves the task correctly (on the grounds that the event *A and B* is in-

cluded in event *A*) must make an effort, consciously or not, to exclude the information described in (ii). This theoretical framework and the claim that representativeness is not sufficient as a theoretical explanation for the Linda problem enables one to explain the effects observed in the study by Wells (1985) that varied the representativeness of the items in Linda-type problems.

Wells (1985) presented subjects with Linda-type problems that employed items containing events that varied in respect to their relevance to the stem. As in the standard problem, his problems contained the three items of interest, represented by events *A, B,* and *A and B*, and asked subjects to give a probability rating for each. The conditions in the experiment differed with respect to the assumed representativeness of the event *A* and the event *B*. One condition employed two representative events, a second employed two nonrepresentative events, and a third employed the sort found on the standard Linda problem, a representative and nonrepresentative event. Now, the standard representativeness hypothesis and the relevance viewpoint make different, in fact opposite, predictions for Linda-type problems containing items of this sort. (It ought to be noted that Wells, 1985, stated that no clear prediction could be made when both components of the conjoined event are representative of the problem's stem.) Assuming that the representativeness heuristic functions additively, it seems clear from the viewpoint of the heuristicist approach that, when both event *A* and event *B* are representative, an item depicting event *A and B* ought to be considered more representative than the item depicting event *A* alone or *B* alone. Therefore, subjects should rate the item depicting the item *A and B* higher than either of the other two items containing event *A* or event *B* alone and commit the error. But from the relevance viewpoint, when each of the two presented items render the stem relevant to a comparable degree, both exploit the information provided in the stem. In such a case, the subject's interpretation of the items need not be misconstrued by considerations of relevance and this leaves subjects relatively unhindered when asked to consider the extensional aspects of the task. And, in fact, Wells (1985) observed an error rate of only 24% (after averaging over four problems) for problems using two events that appear strongly relevant to the stem. This clearly supports the pragmatic prediction based on relevance. He also found an error rate of only 7% on problems (averaging over four problems) that used two events that appear to be nonrepresentative in relation to the stem and found the typical high error rate (72% after averaging over four problems) on problems using one representative and one nonrepresentative event for *A* and *B*. Thus, it appears that the conjunction effect

relies strongly on the anomaly of having two events that differ from each other with regard to their respective relevance to the stem.

Curiously, Kahneman and Tversky (1982) discussed at length the question–answer paradigm in relation to Grice's (1975) maxims. They mentioned that subjects "will generally expect to encounter a cooperative experimenter, although this expectation is often wrong. Because the presentation of irrelevant information violates rules of conversation, subjects are likely to seek relevance in any experimental message" (p. 133). They concluded that "the conversational aspect of judgment studies deserves more careful consideration than it has received in past research, our own included" (p. 135).

Considerations of conversational principles refute a common argument against the possibility of an implicature being attached to the item depicting the nonrepresentative event *A* alone. The argument hinges on the between-subjects manipulation in which one group of subjects has to rate the probability of event *A* alone and another group has to rate the probability of event *A and B*, where *B* is representative of the character described in the stem and *A* is not (Tversky and Kahneman, 1983). Since the results show a higher probability rating for the latter group than for the former in this experiment, in which the *A* item and the *A and B* item simply do not appear together, it is generally concluded that the conjunction effect cannot be due to an implicature. However, this between-subjects manipulation is different from the standard problem and no implicature is needed to explain the effect. In the absence of a direct comparison between the two events, event *A* and event *A and B*, it is impossible to make any claims about subjects' ability to reason extensionally and about the generation of the hypothesized conversational implicature. In this kind of experiment, the subjects' only basis for decision is the degree to which their expectations, determined by the features of the stem, are satisfied by the provided items. Of course, the item depicting event *A and B* meets those expectations better, i.e., is more relevant, in relation to the stem than is the item depicting event *A* alone. Subjects' efforts to modify the item depicting event *A* alone in order to make the task more coherent simply do not come into play in such a demonstration.

In conclusion, putting together the present results with the results from other experiments reported above, one can argue that the way subjects solve the standard problem is not uniform across subjects. The results here allow one to distinguish roughly between the following three main strategies. First, whatever the manipulation used to suppress the error, around one-quarter of the subjects still commit it. These subjects are those for whom the only determinant of the response is the repre-

sentativeness heuristic as explained by the heuristicist approach. Second, according to the present results, around one-half of the subjects interpret the item depicting an event *A* as *A and not-B*. It is argued here that this occurs when the task satisfies the following two criteria: (i) The item format asks subjects to compare an event *A* with an event *A and B*, and (ii) the event that represents the subclass, event *B* in such a comparison, is more relevant to the stem than is the event *A*. It is likely, then, that the two items of concern become event *A and not-B* and event *A and B*, at which point the subject compares each item with the stem and selects the one that fits the stem best, which is *A and B*. To this extent, a representativeness judgment applies, but only after the pragmatic normalization has been performed. This pragmatic normalization is more likely to occur upon an item that has little relevance to the stem, as in the Linda problem's item 1 that ostensibly depicts event *A* alone. Finally, up to one-quarter of the subjects do not commit the error. Those subjects presumably are adept at test taking and problem solving. They are aware of some of the contradictions between the conventions of logic and the rules of language use (Politzer, 1986). By using the conventions of logic, they not only observe the inclusion relationship between events *A* and *A and B*, which enables them to resist the conversational implicature, but they also disregard the pragmatic requirement of a relevance relation between the problem's stem and its demands.

APPENDIX

The English translation of the stem and items to be ranked in the French language equivalent of the Linda problem:

The Helene Problems

Helene is 31 years old. She is single, outspoken and very bright. As a student, she was very involved with issues concerning discrimination against immigrants and concerning social justice, and also participated in antinuclear demonstrations.

( ) Helene is a bank teller.
( ) Helene is active in the feminist movement.
( ) Helene is a bank teller and is active in the feminist movement.

The English translation of the stem and items to be ranked in the problems using novel character descriptions presented in the experiment:

The Lendl Problem

In 1988, Lendl was the best tennis player in the world. Assuming that he will compete in the 1989 French Open Tennis Tournament, consider the following predictions regarding his performance.

Implicit version:
( ) Lendl will play in the finals.
( ) Lendl will win the finals.
( ) Lendl will lose the finals.
Explicit version:
( ) Lendl will play in the finals.
( ) Lendl will play in the finals and win.
( ) Lendl will play in the finals and lose.

The Daniel Problem

In high school, Daniel was always good at math and science. He likes human contact, he has a strong sense for helping others, and he is very determined. Here are some statements regarding Daniel's studies.

Implicit version:
( ) He entered Medical School.
( ) He dropped out of Medical School for lack of interest.
( ) He graduated from Medical School.
Explicit version:
( ) He entered Medical School.
( ) He entered Medical School and dropped out for lack of interest.
( ) He entered Medical School and graduated.

The Edouard Problem

Edouard is a coach. He is 39 years old, courageous and rather hot tempered. He is of the opinion that the police are not doing their job well and that justice is not strict enough. He favors the death penalty.

Here are statements that describe Edouard's reaction in the event that a burglar would threaten him:

Implicit version:
( ) He will not hesitate to use violence against the aggressor.
( ) He will not hesitate to injure the aggressor.

( ) He will not hesitate to kill the aggressor.

Explicit version:

( ) He will not hesitate to use violence against the aggressor.

( ) He will not hesitate to use violence against the aggressor and to injure him.

( ) He will not hesitate to use violence against the aggressor and to kill him.

REFERENCES

Agnoli, F., & Krantz, D. H. (1989). Suppressing natural heuristics by formal instruction: The case of the conjunction fallacy. *Cognitive Psychology, 21*, 515–550.

Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research, 50*, 123–129.

Grice, H. P. (1975). Logic and conversation. In P. Cole, & J. L. Morgan (Eds.), *Syntax and Semantics* (Vol. 3). New York: Academic Press.

Inhelder, B., & Piaget, J. (1964). *The early growth of logic in the child: Classification and seriation.* New York: Norton.

Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. *Cognition, 11*, 123–141.

Kasher, A. (1976). Conversational maxims and rationality. In A. Kasher (Ed.), *Language in focus: Foundations, methods and systems.* Dordrecht, Holland: Reidel.

Leach, C. (1979). *Introduction to statistics: A non-parametric approach for the social sciences.* New York: Wiley.

Leddo, J., Abelson, R. P., & Gross, P. H. (1984). Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology, 47*, 933–943.

Locksley, A., & Stangor, C. (1984). Why versus how often: Causal reasoning and the incidence of judgmental bias. *Journal of Experimental Social Psychology, 20*, 470–483.

McGarrigle, J., Grieve, R., & Hughes, M. (1978). Interpreting inclusion: A contribution to the study of the child's cognitive and linguistic development. *Journal of Experimental Child Psychology, 26*, 528–550.

Morier, D. M., & Borgida, E. (1984). The conjunction fallacy: A task specific phenomenon? *Personality and Social Psychology Bulletin, 10*, 243–252.

Politzer, G. (1986). Laws of language use and formal logic. *Journal of Psycholinguistic Research, 15*, 47–92.

Politzer, G. (1989). *Class-inclusion questions and conversational rules.* Paper presented to the Franco-British Research Group on Philosophy and Cognitive Science. June 1989, Birkbeck College, London.

Shipley, E. F. (1979). The class inclusion task: Question form and distributive comparison. *Journal of Psycholinguistic Research, 8*, 301–331.

Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition.* London: Blackwell.

Tversky, A., & Kahneman, D. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgments under uncertainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*, 293–315.

Wells, G. L. (1985). The conjunction error and the representativeness heuristic. *Social Cognition, 3*, 266–279.

Yates, J. F., & Carlson, B. W. (1986). Conjunction errors: Evidence for multiple judgment procedures, including "signed summation." *Organizational Behavior and Human Decision Processes, 37*, 230–253.