

On the reality of the conjunction fallacy

ASHLEY SIDES, DANIEL OSHERSON, NICOLAO BONINI, and RICCARDO VIALE
Rice University, Houston, Texas

Attributing higher “probability” to a sentence of form p -and- q , relative to p , is a reasoning fallacy only if (1) the word *probability* carries its modern, technical meaning and (2) the sentence p is interpreted as a conjunct of the conjunction p -and- q . Legitimate doubts arise about both conditions in classic demonstrations of the conjunction fallacy. We used betting paradigms and unambiguously conjunctive statements to reduce these sources of ambiguity about conjunctive reasoning. Despite the precautions, conjunction fallacies were as frequent under betting instructions as under standard probability instructions.

The Conjunction Fallacy

Here is the famous Linda story, to be labeled E (for *evidence*) in what follows:

- (E) Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in antinuclear demonstrations.

The task is to rank various statements “by their probability,” including these two:

- (B) Linda is a bank teller.
($B \wedge F$) Linda is a bank teller and is active in the feminist movement.

A majority of respondents across a variety of studies ranked $B \wedge F$ as more probable than B (see Hertwig & Chase, 1998, for a review of findings; the original report is Tversky & Kahneman, 1983). This judgment is in apparent violation of the conjunction law $\Pr(X \wedge Y | Z) \leq \Pr(X | Z)$ for any statements X, Y, Z , with strict inequality for nontrivial cases such as the present example.

The law is not violated, however, if participants in these studies understand the word *probability* in a sense different from the one assigned to it by modern probability theory. There is similarly no violation if B is interpreted to mean $B \wedge \neg F$ or is interpreted in any way other than as a conjunct of $B \wedge F$. The need for clarity about these issues is discussed in the remainder of the present section. We then describe experiments in which we attempted to provide a sharper test of the thesis that naive conjunctive reasoning can be led into fallacy.

Let us first note that we do not attempt to defend naive reasoning by denying the defective character of the judgment $\Pr(X \wedge Y | Z) > \Pr(X | Z)$ (if such a judgment is ever made). In particular, we believe the concept of probability can be sensibly applied to single events (like man reaching

Mars before 2050) and is governed by principles familiar from discussions of Bayesianism (as in Earman, 1992; Horwich, 1982; and Howson & Urbach, 1993). All the events that figured in our experiments were singular in character, resisting placement in classes of similar cases that allow for a meaningful frequency count.

Interpreting the Word Probability

As documented in Hertwig and Gigerenzer (1999), *probability* is polysemous in the general population. It has often been noted, moreover, that, through much of its premodern history, the term *probable* carried a connotation of “approvable opinion” (see Hacking, 1975, chap. 3). Appeal to authority was one way that an opinion was approvable, but another was via evidential support. Thus, John Locke (1671) defined probable propositions as those “for which there be arguments or proofs to make it pass or be received for true” (cited in Krause & Clark, 1993, p. 71). A respondent working with the latter interpretation of probability would attempt to determine whether E provides more support for B or for $B \wedge F$. In what follows, we formalize support in a familiar way and observe that it justifies the intuition that E provides greater support for $B \wedge F$ than for B . Several alternative formalizations would serve our purposes just as well, but we do not attempt a survey of possibilities. Our point is that at least one plausible reading of *probable* exculpates reasoners from the conjunction fallacy.

Many authors agree that a statement X supports a statement Y to the extent that $\Pr(Y | X)$ exceeds $\Pr(Y)$ (see the references cited in Fitelson, 1999, in which the term *confirmation* is used in place of *support*). A simple way to quantify this relation is via the quotient $\Pr(Y | X) / \Pr(Y)$ (the difference works just as well). Here \Pr denotes probability in the modern, technical sense, and the quotient $\Pr(Y | X) / \Pr(Y)$ translates the support concept into modern terms. According to the definition, E supports $B \wedge F$ more than E supports B if and only if

$$\frac{\Pr(B \wedge F | E)}{\Pr(B \wedge F)} > \frac{\Pr(B | E)}{\Pr(B)}$$

The authors thank Andrea Cerroni, Karin Dudziak, Denise Wu, Andrea Pozzali, and Zhihua Tang for assistance in performing Experiment 1. Correspondence should be addressed to D. Osherson, Department of Psychology, MS-25, P. O. Box 1892, Rice University, Houston, TX 77251-1892 (e-mail: osherson@rice.edu).

An application of Bayes's theorem reveals that the inequality above holds if and only if $\Pr(E|B \wedge F) > \Pr(E|B)$. And the latter inequality is entirely reasonable since it asserts that Linda is more likely to be single, outspoken, and so on, on the assumption that she is a feminist bank teller than on the mere assumption that she is a bank teller. Hence, one interpretation of majority choices in the conjunction problem is that (1) most respondents have a support interpretation of *probability*, (2) their conception of the support that statement X provides for Y can be formalized as the ratio of $\Pr(Y|X)$ to $\Pr(Y)$, and (3) they accurately perceive E to provide more support for $B \wedge F$ than for B , in the foregoing sense of support.

The hypothesis that the word *probability* is interpreted as support is not the same as claiming that the standard probability of B given E is the probability of E given B , and similarly for $B \wedge F$. The latter claim is explored in Wolford, Taylor, and Beck (1990) and aptly criticized by Bar-Hillel (1991; see also Wolford, 1991). On the other hand, the support hypothesis is close to an analysis advanced in Hertwig and Chase (1998). The latter discussion is complicated, however, by a relatively indirect measure of support (due to Nozick, 1981).

If responses to the Linda problem are often due to a support interpretation of *probability*, we would expect more conformity to the conjunction law when the wording of the problem discourages this interpretation. One such wording is due to Fiedler (1988), who asked participants to estimate "how many out of 100 people who are like Linda" satisfy B and $B \wedge F$. Relative to the original problem, it seems harder to construe Fiedler's frequency question as involving support. In fact, conjunction violations were substantially depressed in response to the frequency wording.¹

Improvement with frequency formats, however, leaves open the possibility of irrational judgment when confronted with single events. To assess conformity to the conjunction law in these circumstances, it seems necessary to frame questions that avoid the word *probability* and its cognates. The obvious strategy is to offer a choice between betting on a conjunctive proposition versus one of its conjuncts. Such was the idea behind Tversky and Kahneman's (1983) problem involving a regular six-sided die with four green faces and two red faces. Participants were asked to select one sequence, from a set of three, with a \$25 prize if the chosen sequence appeared embedded within 20 rolls of the die. A majority of participants preferred to bet on *GRGRRR* than on *RGRRR*, which violates the conjunction law, since the former sequence includes the latter. In this case, however, the violation might have resulted from failure to notice the inclusion, which is not perceptually salient. Indeed, in a further condition, a majority of students preferred a probabilistically correct argument for choosing the included sequence over an argument based on the proportions of red and green outcomes. In the preferred argument, the inclusion relation was made explicit. Tversky and Kahneman also tested a version of the Linda problem in which participants were asked which

of B and $B \wedge F$ they preferred to bet on. There was some decline in the rate of conjunction violation, but it nonetheless characterized a majority of responses.

The betting version of the Linda problem strikes us as inconclusive, however, in light of the hypothetical character of the question. Linda is not a real person, and no bets will be paid. Respondents may consequently be inclined to interpret the query as a disguised probability question, leading to the same ambiguity as before. Similar concerns beset the betting questions posed in Wolford et al. (1990). Genuine bets were made (and paid) in Bar-Hillel and Neter's (1993) careful study of violations of the disjunction law (according to which the probability of a given event cannot be higher than the probability of any event that includes it). But bets in their experiments concerned people with whom the participants were unfamiliar and outcomes already known to the experimenter. This is not the usual kind of betting context. Wagers on sporting events, for example, bear on familiar teams and as-yet-undetermined outcomes. In the betting experiments reported below, we therefore prepared questions involving future events about which our participants possessed background knowledge (the latter plays the role of evidence E).

Betting paradigms also discourage interpreting the probability of a sentence as its expected informational value. To explain, suppose that people judge $B \wedge F$ to be much more informative than B , but only slightly less probable. Then the product of probability and informativeness (the *expected informativeness*) might be higher for the conjunction than for its conjunct. In response to a question about probability, participants might choose the alternative with higher expected informativeness on conversational grounds (since one goal of polite conversation is to be maximally informative). This possibility was recognized by Tversky and Kahneman (1983) and was also discussed by Bar-Hillel and Neter (1993). Since conversational goals are clearly irrelevant to gambling, betting paradigms are well suited to minimize the impact of expected informativeness on respondents' choices.

Ambiguous Logical Form

Answers to the Linda problem constitute a conjunction fallacy only if the options labeled $B \wedge F$ and B are interpreted as a conjunction and one of its conjuncts. It has been widely observed that, in the presence of the alternative $B \wedge F$, the pragmatics of conversation may lead respondents to interpret B as $B \wedge \neg F$ —that is, to interpret "Linda is a bank teller" as "Linda is a bank teller and not a feminist" (Dulany & Hilton, 1991; Macdonald & Gilhooly, 1986; Morier & Borgida, 1984; Politzer & Noveck, 1991). Attempts to prevent such an interpretation began with Tversky and Kahneman (1983), who reported conjunction errors even when B was rendered as "Linda is a bank teller whether or not she is active in the feminist movement." As noted in Hilton (1995, p. 260), however, such a formulation might be interpreted as asserting that Linda is a bank teller even if she is a feminist. Whatever the

proper interpretation of the latter assertion, it seems different from plain B . An attempt to disambiguate conjunctions through presupposition is reported in Politzer and Noveck (1991), but substantial conjunction errors persist.

When different groups of respondents assign probabilities to B versus $B \wedge F$, higher numbers typically are attached to the latter (Tversky & Kahneman, 1983). Pragmatic explanation of such results is obviously ruled out. Although higher average probabilities for conjunctions, relative to conjuncts, suggest a tendency to violate the conjunction principle, no individual actually commits a fallacy (since no individual assigns probabilities to both B and $B \wedge F$). The separate-groups paradigm is therefore unsuited to assess the prevalence of conjunction mistakes.

In addition to pragmatic influences on the interpretation of B , it is possible that $B \wedge F$ is not coded by participants as a conjunction. Mellers, Hertwig, and Kahneman (2001) discuss the possibility that expressions such as “bank teller and feminist” are interpreted disjunctively, which would justify answers that have been considered fallacious heretofore. Supporting this conjecture are sentences such as “We invited friends and colleagues to the party,” in which the class of invitees can be read as embracing anyone who is either a friend or a colleague (or both). The authors describe data that contradict this explanation, however, inasmuch as they document numerous conjunction fallacies in a version of the Linda problem that relies on the conjunctive expression “feminist bank teller” (and poses the question in frequency format).

Ambiguity surrounding the word *and* seems not to arise when it is used to connect propositions rather than categories. Thus, the sentence “IBM stock will rise tomorrow and Disney stock will fall tomorrow” seems immune to a disjunctive reading. For this reason, the stimulus materials figuring in the experiments reported below involved only conjoined sentences instead of conjoined noun phrases.

In summary, considerable research has been devoted to the question, Does naive reasoning about probability respect the conjunction principle? Previous studies have provided valuable information, but the matter remains open because of (1) potential ambiguity surrounding the word *probability*, (2) concerns about misleading pragmatic influences, and (3) uncertainty about the conjunctive reading of categories united by the word *and*. The experiments reported below attempted to address these issues.

EXPERIMENT 1

In Experiments 1–3, we attempted to clarify the conjunction problem in the following ways. First, betting instructions were used for half of the participants, and standard probability instructions were used for the other half. Second, only two alternatives appeared in each item (in contrast to the original version of the Linda problem, which employed multiple alternatives; Tversky & Kahneman, 1983). Third, all the items involved future events whose occurrence was equally unknown to both subject and experi-

menter. Fourth, conjunctions involved sentences rather than noun phrases, discouraging disjunctive interpretation.

In the instructions for the betting condition, it was emphasized that payoffs would be made if the selected event occurred, regardless of whether the unselected event occurred. In this way, we hoped to make clear that both $X \wedge Y$ and X could be true, which implies that X does not mean $X \wedge \neg Y$. A more direct means of preventing unwanted conversational inferences is described in Experiment 4.

Method

Stimuli. Twenty-four pairs of events were constructed, including those shown in Table 1. Nine pairs had the logical forms $X, Y \wedge X$. Nine additional pairs had the forms $X, Y \wedge Z$, and another 6 pairs had the forms X, Y . The 9 pairs of the form $X, Y \wedge X$ were critical to assessing the conjunction error. They will be called *fallacy items* in what follows. The remaining 15 pairs were used as fillers.²

All the events involved people, places, or circumstances familiar to college students in Houston. Whether the event would occur depended on future outcomes that could be objectively verified at a date determined by the wording of the event. In the fallacy items, the conjunct Y was designed to enhance the plausibility of X (i.e., $\Pr(X|Y) > \Pr(X)$, in the authors' opinion).

The 24 (9 + 9 + 6) event pairs were used to construct two sets of booklets. In the *probability* booklets, each page displayed the two events along with the instruction to mark the one “most likely to occur.” The *betting* booklets were the same except that the instruction was to mark the event “you would like to bet on.” The order in which the two events appeared on a page was chosen randomly. The order of the pages in a booklet was individually randomized, with the constraint that two fallacy items never occur consecutively.

Participants. The participants were 90 undergraduates from Rice University who were fulfilling course requirements in the autumn of 1998. They were run in sessions of 12–40 students, depending on scheduling vicissitudes. In each session, the participants were randomly assigned to a probability group or a betting group. In all, 45 students participated in the probability condition, and 45 participated in the betting condition.

Procedure. Each participant in the probability condition received a probability booklet. They were asked to consider each page individually and to mark the event (out of two) of “highest probability” (i.e., the one “most likely to occur”).

In the betting condition, a \$50 bill was prominently displayed at the start of the session. The participants were then informed that the experimental procedure had two parts. First, they would mark one event on each page of their booklets—namely, the event on which they preferred to bet. Then, one student would be designated at random, and one page from his/her booklet would be randomly selected. If the event marked on that page occurred by the relevant date, the participant would be awarded the \$50 prize (students supplied electronic and postal addresses for collection purposes). A new \$50 prize applied to each of the betting sessions, irrespective of the number of participants (which ranged from 6 to 20). In the betting condition, reference to expressions, such as *probability*, *chance*, and *likelihood*, was scrupulously avoided. In the course of the instructions, it was repeated several times that if the event marked on the chosen page occurred by the relevant date—irrespective of the unselected event—then the chosen participant would win \$50.

Results

In the probability condition, 38 of 45 participants committed the conjunction error at least once (i.e., labeling a conjunction as more likely than its conjunct).

Table 1
Two Examples of Each Kind of Event Pairs Used in Experiments 1 and 2

X	Y \wedge X
1. The percentage of adolescent smokers in Texas will decrease at least 15% from current levels by September 1, 1999.	The cigarette tax in Texas will increase by \$1.00 per pack and the percentage of adolescent smokers in Texas will decrease at least 15% from current levels by September 1, 1999.
2. By September 1, 1999, an experimental vaccine for childhood leukemia will be announced.	The National Institutes of Health (NIH) will increase spending on vaccine development by 50% in the first 9 months of 1999, and by September 1, 1999, an experimental vaccine for childhood leukemia will be announced.
X	Y \wedge Z
3. The University of Houston Philosophy Department will hire 3 new faculty members by September 1, 1999.	The University of Houston writing department will be rated in the top 10% nationwide and will announce that it will expand its faculty by September 1, 1999.
4. By September 1, 1999, Texas will require people to pass a literacy test before serving on a jury.	By September 1, 1999, Texas will start selecting juries from a pool of licensed drivers rather than registered voters and the number of registered voters will increase by 10%.
X	Y
5. Bill Clinton will announce his intention to seek a divorce before September 1, 1999.	By September 1, 1999, Janet Reno will announce her intention to run for the Presidency.
6. Fidel Castro will be removed from political power in Cuba by September 1, 1999.	U.S. forces will be sent to Havana, Cuba before September 1, 1999.

Out of 9 possible occasions, the average number of errors per participant was 3.4 ($SD = 2.55$). In the betting condition, 36 of 45 participants committed at least one conjunction error (choosing to bet on a conjunction rather than its conjunct), with an average of 3.2 errors ($SD = 2.33$). A *t* test revealed that the difference in means did not approach significance. Table 2 shows the number of participants in each condition who made *m* errors, for *m* between 0 and 9.

For a given fallacy item and a given condition, call the proportion of participants who committed the conjunction error the *fallacy score* for that item in that condition. We correlated the fallacy scores for the 9 items across the two conditions. The Pearson coefficient was .82 ($p < .01$), suggesting similar mental processes in the two conditions.

No fallacy was committed, of course, if the participants' responses resulted from inattention or lack of interest. Inattentive responding would favor equal fallacy rates across the nine different fallacy items. But, in fact, the 9 items attracted very different numbers of fallacy responses in both the betting condition and the probability condition. For example, combining across the two conditions, Item 1 in Table 1 attracted 51 fallacy responses (out of 90 possible), whereas Item 2 attracted 19 fallacy responses. Thirty-four participants committed the fallacy for Item 1 but not for Item 2, whereas only 2 participants had the reverse profile. This difference was reliable by a binomial test in which inattentive responding was assimilated to the toss of a fair coin ($p < .05$, two-tailed). The same test yielded a reliable difference between Item 1 and every other item except for two. We conclude that our participants' fallacious responding was not due to inattention to the task.

EXPERIMENT 2

To test the robustness of our findings, Experiment 1 was replicated with a new group of students at a different university.

Method

Stimuli. With a few exceptions, the stimuli from Experiment 1 were used for Experiment 2. The exceptions arose from events transpiring during the interval separating the experiments. For example, the event "the new *Star Wars* movie will receive two thumbs down from Siskel and Ebert by September 1, 1999" needed replacement after the death of Gene Siskel early that year.

Participants. The participants in Experiment 2 were 57 undergraduate volunteers from the University of Houston, a public institution with diverse enrollment, located in downtown Houston.

Procedure. The procedure was the same as described above for Experiment 1. Twenty-nine participants were randomly assigned to the probability condition, and 28 participants were randomly assigned to the betting condition. The participants were run in groups of 3–14 in the first months of 1999.

Results

In the probability condition, all 29 participants committed at least one conjunction error, with an average of 5.93

Table 2
Number of Participants in the Probability and Betting Conditions of Experiment 1 Who Committed *m* Conjunction Errors ($0 \leq m \leq 9$)

Condition	Number of Conjunction Errors									
	0	1	2	3	4	5	6	7	8	9
Probability	7	8	1	6	8	5	4	3	2	1
Betting	9	5	3	6	7	8	5	1	0	1

Note—For each condition, $n = 45$ participants.

Table 3
Number of Participants in the Probability and Betting
Conditions of Experiment 2 Who Committed *m*
Conjunction Errors ($0 \leq m \leq 9$)

Condition	Number of Conjunction Errors									
	0	1	2	3	4	5	6	7	8	9
Probability	0	0	0	2	3	8	5	5	5	1
Betting	3	2	1	2	1	5	7	3	3	1

Note—For the probability and betting conditions, *ns* = 29 and 28 participants, respectively.

out of 9 possible errors (*SD* = 1.60). In the betting condition, 25 of 28 participants committed at least one conjunction error, with an average of 4.86 errors (*SD* = 2.65). The difference in means was marginally significant [$t(55) = 1.86, p < .07$]. Table 3 shows the number of participants in each condition who made *m* errors, for *m* between 0 and 9. The Pearson coefficient for the correlation of fallacy scores for the 9 fallacy items across the probability versus betting conditions was .84 ($p < .01$).

As in Experiment 1, the 9 items attracted very different numbers of fallacy responses in both the betting condition and the probability condition. Combining across conditions, the item with top rank in attracting fallacious responses was reliably more attractive than every other item except three (the binomial test was used for each comparison, as in Experiment 1).

A univariate analysis of variance revealed significantly more conjunction fallacies committed by the University of Houston students than by the Rice University students ($F = 27.6, p < .001$), but no interaction with experimental condition. The scholastic preparation of Rice University students is superior to that of University of Houston students according to such measures as SAT scores and high school GPA. The lower fallacy rate among the former, relative to the latter, was thus consistent with Stanovich (1999), who reported higher mean SAT scores for students who avoided fallacy on conjunction problems.

Combining the results of Experiments 1 and 2 above, we observed an average of 4.42 conjunction errors in the probability condition ($n = 74, SD = 2.53$) and an average of 3.82 conjunction errors in the betting condition ($n = 73, SD = 2.57$). The difference did not reach statistical significance.

EXPERIMENT 3

As a further test of robustness, we replicated Experiments 1 and 2 using Italian participants and items adapted to Italian culture. We attempted to conserve the broad character of the American stimuli through small changes in names and places. For example, reference to the National Institutes of Health was replaced by “L’Istituto Superiore di Sanità.”

Method

Ninety-two students in a sociology class at the University of Milan performed the probability version of the procedure. Sixty ad-

ditional students in a different sociology class performed the betting version.

Results

In the probability condition, 90 of 92 participants committed at least one conjunction error, with an average of 4.33 of 9 possible errors (*SD* = 1.61). In the betting condition, all 60 participants committed at least one conjunction error, with an average of 4.85 errors (*SD* = 1.75). The difference in means is marginally significant by *t* test [$t(150) = 1.88, p < .06$]. Table 4 shows the number of participants in each condition who made *m* errors out of 9. The Pearson coefficient for the correlation of fallacy scores for the 9 fallacy items across the two conditions was .96 ($p < .01$). Combining across the two conditions, the item with top rank in attracting fallacious responses was reliably more attractive than every other item except for one (using the binomial test, as in Experiment 1).

The results of Experiments 1–3 support the thesis that conjunction fallacies are a genuine feature of naive reasoning about chance. The rate of conjunction errors was high under betting instructions in all three experiments, with only 9% of the betting participants avoiding the fallacy altogether. Moreover, each critical problem involved just the options *X* and $X \wedge Y$ accompanied by the remark that payoffs were contingent on the occurrence of the chosen event irrespective of the status of the other one. In these options, *X* and *Y* were sentences rather than noun phrases, which rendered a nonconjunctive interpretation unlikely.

EXPERIMENT 4

It remains possible that, in Experiments 1–3, the participants interpreted conjuncts *X* as implicit conjunctions $X \wedge \neg Y$, thus misinterpreting the question for pragmatic reasons. The present experiment attempted to rule out this possibility.

Suppose that a third party will be shown which of two sentences you choose, and not the other one. Then, this person’s interpretation of your chosen sentence cannot depend on the unselected sentence. The procedure of the present experiment exploited this fact to dissuade the participants from interpreting *X* as $X \wedge \neg Y$ in the context of $X \wedge Y$. The participants were told that a judge who sees only one of the two sentences (i.e., the one selected by the par-

Table 4
Number of Participants in the Probability and Betting
Conditions of Experiment 3 Who Committed *m*
Conjunction Errors ($0 \leq m \leq 9$)

Condition	Number of Conjunction Errors									
	0	1	2	3	4	5	6	7	8	9
Probability	2	3	8	8	30	19	14	8	0	0
Betting	0	2	2	11	9	15	12	4	4	1

Note—For the probability and betting conditions, *ns* = 92 and 60 participants, respectively.

ticipant) would determine whether the sentence came true and pay off bets accordingly.

In addition, half the participants confronted conjunctions involving the particle *and*, whereas the other half saw parallel items involving other conjunctive forms (e.g., *after which*). In this way, it could be seen whether errors depend narrowly on the manner in which conjunctions are expressed in English.

Method

Stimuli. Ten pairs of events were constructed. Two pairs (the fallacy items) had the logical forms $X, Y \wedge X$. Eight additional pairs were used as fillers. Each conjunction was rendered either *explicitly* using *and* or *implicitly* via another construction. Table 5 shows the contrast between explicit and implicit conjunctions, including all four fallacy items.

The 10 event pairs were used to construct two sets of booklets, as in Experiments 1–3. In the explicit booklets, all conjunctions were rendered explicitly. In the implicit booklets, they were rendered implicitly. The order in which the two events appeared on a page was chosen randomly. The order of the pages in a booklet was individually randomized, with the constraint that the two fallacy items not occur consecutively. At the top of each page of each booklet was the following instruction (amplified by the oral instructions):

Please choose the sentence on this page that you would prefer to bet on, and scratch out the other sentence. An independent judge who knows nothing of this experiment will determine which bets will be paid off (50 cents per question) based only on the sentence left legible on this page.

Participants. Eighty-three undergraduates from Rice University participated for partial fulfillment of course requirements. They were randomly assigned to either the explicit condition ($n = 43$) or the implicit condition ($n = 40$).

Procedure. The participants were asked to select one event on each page of their booklets and to mark out the event on which they did not want to bet. They were instructed to make that sentence illegible, so that an independent judge could not tell what event had been deleted. They were told that such a judge with no knowledge of the experiment had been commissioned to decide which bets to pay off when all bets came due. The judge would make her decision entirely on the basis of the event that she could read on each page (more than one legible event on a page would disqualify both). It was emphasized that the judge would not have seen any of the items beforehand and would make her decisions solely by examining the legible sentences remaining in the participant's booklet.

The participants were told that each page of their booklets was potentially worth 50 cents. Prior to making their choices they were presented with a claim ticket that corresponded to the identification number on their booklets and were instructed to return to the experimenter's office on a specified date 4 months after the experiment took place, when the foretold events had either transpired or not. (Bets were duly paid off on the appointed day.)

Results

In the explicit condition, 28 of the 43 participants committed the conjunction fallacy at least once (out of two possible occasions); 7 committed it both times. In the implicit condition, 25 of the 40 participants committed the conjunction fallacy at least once; 7 committed it both times. There was no significant difference between the fallacy rates in the two conditions, so they were combined to determine whether one item was more conducive to fallacy than the other. (In what follows, Item 3 in Table 5 is assimilated to Item 1, and Item 4 is assimilated to Item 2.) Item 1 attracted 21 fallacies out of 83 participants, and

Table 5
Some Items Used in Experiment 4

Item Type	First Event	Second Event
1. Explicit Fallacy	A tax cut will be passed by Congress between January 1 and March 31, 2000.	A tax cut will be passed by Congress between January 1 and March 31, 2000 and it will be supported by most Democrats.
2. Explicit Fallacy	Parents in Houston will be required to attend a "good sportsmanship" workshop before their children are allowed to participate in organized sports.	Several incidents will take place involving parental conflict at little league games, and parents in Houston will be required to attend a "good sportsmanship" workshop before their children are allowed to participate in organized sports.
3. Implicit Fallacy	A tax cut will be passed by Congress between January 1 and March 31, 2000.	A tax cut will be passed by Congress between January 1 and March 31, 2000 with the support of most Democrats.
4. Implicit Fallacy	Parents in Houston will be required to attend a "good sportsmanship" workshop before their children are allowed to participate in organized sports.	Several incidents will take place involving parental conflict at little league games, after which parents in Houston will be required to attend a "good sportsmanship" workshop before their children are allowed to participate in organized sports.
5. Explicit Filler	A budget freeze that will prevent salary increases for Texas public school teachers will be implemented, and Texas public school teachers will go on strike.	Another instance of guns brought to school by students in Texas will occur, and Texas public school teachers will go on strike until metal detectors are installed in schools.
6. Implicit Filler	A budget freeze that will prevent salary increases for Texas public school teachers will be implemented, then Texas public school teachers will go on strike.	Another instance of guns brought to school by students in Texas will occur, then Texas public school teachers will go on strike until metal detectors are installed in schools.

Item 2 attracted 46 fallacies out of 83 participants. Thirty-nine participants committed the conjunction fallacy on just one of the two items. Of this group, 32 participants committed it for Item 2, and only 7 participants committed it for Item 1.

Experiment 4 induced numerous violations of the conjunction principle. Choice among bets was the dependent variable with no mention of *probability*, so the fallacy is not likely to result from nonmodern interpretation of the probability idiom. Moreover, the instruction to leave only one choice legible for an independent judge can be expected to cancel any pragmatic tendency to interpret X as $X \wedge \neg Y$. Finally, the fallacy rate was virtually the same in the explicit and implicit conditions, suggesting that the conjunction fallacy does not depend narrowly on use of *and* to express conjunction.

DISCUSSION

Faced with evidence of fallacious reasoning about single events, some psychologists deny that there is anything counternormative in attributing greater chance to a conjunction of such events than to one of the conjuncts (Gigerenzer, 1991; Gigerenzer, Hoffrage, & Kleinbolting, 1991). Putting aside semantic issues about the word *chance*, however, it is surely irrational to prefer a reward in case an event of form $X \wedge Y$ occurs, relative to receiving the same reward in case X occurs (as observed in Kahneman & Tversky, 1996). And just this pattern occurred throughout our four experiments.

Another defense of naive judgment is the claim that subjects often understand X as $X \wedge \neg Y$ in the context of $X \wedge Y$ (Hilton, 1995). This interpretation is said to be consistent with a cooperative attitude about discourse (Grice, 1975). Previous attempts to control for such variables (as in Tversky & Kahneman, 1983) have been deemed insufficient. The procedure of Experiment 4 will perhaps prove more persuasive since it invokes a judge who must interpret X outside of the frame provided by $X \wedge Y$. Indeed, inability to interpret X as X in these conditions would cast doubt on people's mastery of the very pragmatic principles used to defend their reasoning about probability.

The general issue brought into focus by the conjunction problem is the extent to which reasoning about chance is sensitive to the constraints imposed by the logical structure of events. Insensitivity leads to *probabilistic incoherence*, and has been documented for events with various logical structures (e.g., Dawes, Mirels, Gold, & Donahue, 1993; Osherson, Lane, Hartley, & Batsell, 2001). An easy demonstration is obtained by asking one's colleagues whether it is sensible to assign 30% probability to man reaching Mars by 2020, 80% probability to a sustained economic downturn before 2010, and 5% probability to the conjunction of the two events. In our experience, most people find nothing objectionable to these estimates, yet they are probabilistically incoherent (see Neapolitan, 1990, p. 128). Note that this case seems not to be open to pragmatic interpretation.

Maintaining coherence is computationally difficult in the sense of requiring resources that grow rapidly as a function of the structural complexity of the judgments in play (see Georgakopoulos, Kavvadias, & Papadimitriou, 1988). The noteworthy feature of human judgment, however, is that incoherence strikes even in simple cases, such as comparing the chances of $X \wedge Y$ and X . Why are we so disinclined to coordinate probability with logical structure?

One reason might be the separation of neural loci for the evaluation of logical implication versus empirical plausibility. These two mental activities induce metabolic activity in opposite hemispheres connected by axon bundles of relatively narrow bandwidth (for neuroimaging evidence, see Osherson et al., 1998, and Parsons & Osherson, 2001). Coordinating assessments of chance with the logical forms of statements may not be facilitated by this cognitive architecture.

REFERENCES

- BAR-HILLEL, M. (1991). Commentary on Wolford, Taylor, and Beck: The conjunction fallacy? *Memory & Cognition*, **19**, 412-414.
- BAR-HILLEL, M., & NETER, E. (1993). How alike it is versus how likely it is: A disjunction fallacy in probability judgments. *Journal of Personality & Social Psychology*, **65**, 1119-1131.
- DAWES, R., MIRELS, H. L., GOLD, E., & DONAHUE, E. (1993). Equating inverse probabilities in implicit personality judgments. *Psychological Science*, **4**, 396-400.
- DULANY, D. E., & HILTON, D. J. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition*, **9**, 85-110.
- EARMAN, J. (1992). *Bayes or Bust?* Cambridge, MA: MIT Press.
- FIEDLER, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research*, **50**, 123-129.
- FITELSON, B. (1999). The plurality of Bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science*, **66**, S362-S378.
- GEORGAKOPOULOS, G., KAVVADIAS, D., & PAPADIMITRIOU, C. (1988). Probabilistic satisfiability. *Journal of Complexity*, **4**, 1-11.
- GIGERENZER, G. (1991). How to make cognitive illusions disappear: Beyond "heuristics and biases." In W. Storche & M. Hewstone, (Eds.), *European review of social psychology* (Vol. 2, pp. 83-115). New York: Wiley.
- GIGERENZER, G. (1996). Reply to Tversky and Kahneman. *Psychological Review*, **103**, 592-593.
- GIGERENZER, G., HOFFRAGE, U., & KLEINBOLTING, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, **98**, 506-528.
- GRICE, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41-58). New York: Academic Press.
- HACKING, I. (1975). *The emergence of probability*. Cambridge: Cambridge University Press.
- HERTWIG, R., & CHASE, V. M. (1998). Many reasons or just one: How response mode affects reasoning in the conjunction problem. *Thinking & Reasoning*, **4**, 319-352.
- HERTWIG, R., & GIGERENZER, G. (1999). The "conjunction fallacy" revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making*, **12**, 275-305.
- HILTON, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin*, **118**, 248-271.
- HORWICH, P. (1982). *Probability and evidence*. Cambridge: Cambridge University Press.
- HOWSON, C., & URBACH, P. (1993). *Scientific reasoning: The Bayesian approach*. Chicago: Open Court.

- JONES, S. K., JONES, K. T., & FRISCH, D. (1995). Biases of probability assessment: A comparison of frequency and single-case judgments. *Organizational Behavior & Human Decision Processes*, **61**, 109-122.
- KAHNEMAN, D., & TVERSKY, A. (1996). On the reality of cognitive illusions. *Psychological Review*, **103**, 582-591.
- KRAUSE, P., & CLARK, D. (1993). *Representing uncertain knowledge*. Norwell, MA: Kluwer.
- MACDONALD, R., & GILHOOLY, K. (1986). More about Linda or conjunctions in context. *European Journal of Cognitive Psychology*, **2**, 57-70.
- MELLERS, B., HERTWIG, R., & KAHNEMAN, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, **12**, 269-275.
- MORIER, D., & BORGIDA, E. (1984). The conjunction fallacy: A task specific phenomenon? *Personality & Social Psychology Bulletin*, **10**, 243-252.
- NEAPOLITAN, R. (1990). *Probabilistic reasoning in expert systems: Theory and algorithms*. New York: Wiley.
- NOZICK, R. (1981). *Philosophical explanations*. Oxford: Oxford University Press, Clarendon Press.
- OSHERSON, D., LANE, D., HARTLEY, P., & BATSELL, R. (2001). Coherent probability from incoherent judgment. *Journal of Experimental Psychology: Applied*, **70**, 3-12.
- OSHERSON, D., PERANI, D., CAPPAS, S., SCHNUR, T., GRASSI, F., & FAZIO, F. (1998). Distinct brain loci in deductive versus probabilistic reasoning. *Neuropsychologia*, **36**, 369-376.
- PARSONS, L., & OSHERSON, D. (2001). New evidence for distinct right and left brain systems for deductive vs. probabilistic reasoning. *Cerebral Cortex*, **11**, 954-965.
- POLITZER, G., & NOVECK, I. A. (1991). Are conjunction rule violations the result of conversational rule violations? *Journal of Psycholinguistic Research*, **20**, 83-103.
- STANOVICH, K. E. (1999). *Who is rational?* Mahwah, NJ: Erlbaum.
- TVERSKY, A., & KAHNEMAN, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, **90**, 293-315.
- WOLFORD, G. (1991). The conjunction fallacy? A reply to Bar-Hillel. *Memory & Cognition*, **19**, 415-417.
- WOLFORD, G., TAYLOR, H. A., & BECK, J. R. (1990). The conjunction fallacy? *Memory & Cognition*, **18**, 47-53.

NOTES

1. But see Jones, Jones, and Frisch (1995), who document a smaller improvement for frequency wording. The observation that frequency formats increase conformity to the conjunction principle is due to Tversky and Kahneman (1983). For divergent interpretations of the increase, see Kahneman and Tversky (1996), Gigerenzer (1996), and references cited there. Hertwig and Chase (1998) make a convincing case that part of the increase is due to the use of numerical responses in frequency versions of the problem, relative to rank-order responses in the original.
2. Some of the items "expired" before all the data could be collected in the sense that an event mentioned in one choice either occurred or became impossible. Items that expired during testing were replaced by similar ones.

(Manuscript received April 27, 2001;
revision accepted for publication November 12, 2001.)