

The conjunction fallacy?

GEORGE WOLFORD and HOLLY A. TAYLOR
Dartmouth College, Hanover, New Hampshire

and

J. ROBERT BECK
Dartmouth Medical School, Hanover, New Hampshire

Tversky and Kahneman (1983) showed that when subjects are asked to rate the likelihood of several alternatives, including single and joint events, they often make a "conjunction fallacy." That is, they rate the conjunction of two events as being more likely than one of the constituent events. This, they claim, is a fallacy, since the conjunction of two events can never be more probable than either of the component events. In addition, they found that prior training in probability theory does not decrease the likelihood of making this fallacy. We argue that in some contexts, an alternative that contains the conjunction of two events can be more probable than an alternative that contains only one of the conjunction's constituent events. We carried out four experiments in which we manipulated this context. The frequency of making a conjunction fallacy was affected by the manipulation of context. Furthermore, when the context was clearly specified, prior training in statistics influenced the ratings.

Many decisions involve some uncertainty. We often do not know the precise outcome that will follow if we select a particular alternative. Will the cancer be eradicated if we select surgery? Will the teachers' union agree to the contract if we offer a 9% raise? Classic decision models suggest that we should specify the outcomes that might follow if an alternative is selected, estimate the probability of each outcome, and determine the value of that outcome. Finally, we should combine those probabilities and values in some meaningful way. Although people may follow these steps, at least informally, in arriving at a decision, they often perform less than optimally at the various stages. Tversky and Kahneman (1982, 1983) have shown that people use a number of heuristics and exhibit several biases in estimating and combining probabilities and values. The description of these biases and heuristics has contributed to our understanding of the way people actually make decisions.

In this article, we concentrate on one of the errors exhibited by people in estimating probabilities, namely, the "conjunction fallacy" (Tversky & Kahneman, 1982, 1983). In a variety of circumstances, subjects rate the conjunction of two events as more probable than one of the constituent events. This is generally considered a fallacy since, in probability theory, the probability of the intersection of two events can never be greater than the probability of either event alone. Tversky and Kahneman (1983) refer to the extensional property to contrast joint

and constituent events. The constituent event should always be favored since, by definition, it is included in the joint event. We plan to show that in some contexts, a different type of judgment is required, and it might be appropriate to rate an alternative containing a joint event as being more likely than an alternative containing one of the constituent events.

The conjunction fallacy usually arises when prior information indicates that some event, *A*, is quite probable and some event, *B*, is quite improbable. In that situation, subjects often rate the intersection of conjunction of Events *A* and *B* as more probable than Event *B* alone.

A common example of the conjunction fallacy is the "Linda problem" (or sometimes the "Bill problem"). Linda is described in the following paragraph from Tversky and Kahneman (1982, 1983):

Linda Problem

Linda is 31, single, outspoken, and very bright. She majored in philosophy in college. As a student, she was deeply concerned with discrimination and other social issues, and participated in anti-nuclear demonstrations.

Subjects are then asked to rank the likelihood of various alternatives, such as: (1) Linda is active in the feminist movement. (2) Linda is a bank teller. (3) Linda is a bank teller and is active in the feminist movement.

Tversky and Kahneman found that between 85% and 90% of their subjects rated Alternative 3, the conjunction, as more likely than Alternative 2, although nearly all of the subjects rated Alternative 1 as the most likely. This result has been replicated using a variety of scenarios and procedural alterations. Using this direct version of the Linda problem, the result appears to be unaffected by

We are grateful for comments on the models by Laurie Snell and Peter Doyle, and for comments on the manuscript by Carol Fowler. Requests for reprints may be addressed to the first author at the Department of Psychology, Dartmouth College, Hanover, NH 03755.

—Accepted by previous editor, Alice F. Healy

prior training in probability and statistics. Subjects with several courses in probability and statistics are as likely to make the fallacy as are untrained subjects.

Tversky and Kahneman argue that subjects make the conjunction fallacy because they use a "representativeness" heuristic to arrive at their judgment. According to this heuristic, every event is assigned some value of representativeness for the object, or person, in question. Being active in the feminist movement is representative of someone with Linda's background. Being a bank teller is not. One major difference between values of representativeness and probability is the way they are combined. The representativeness of a combination in some cases could be a linear combination of the representativeness of the component characteristics. The representativeness of a combination, then, could be higher than the component with the lowest value of representativeness. Tversky and Kahneman measured the representativeness of the various alternatives in the Linda problem and showed that the observed likelihood rankings were consistent with the representativeness ratings.

Other explanations for the conjunction fallacy have been considered. Several investigators have explored the possibility that subjects fail to understand the various alternatives (Locksley & Stangor, 1984; Morier & Borgida, 1984; Tversky & Kahneman, 1983). In particular, it seemed possible that subjects might misinterpret the single-event alternatives in the Linda problem to imply the negation of the other event. For instance, subjects might read Alternative 2, "bank teller," to imply that Linda was a bank teller and was *not* active in the feminist movement. Under this interpretation, that alternative might well be less likely than the joint event in Alternative 3. These investigators clarified the meaning of the various alternatives by adding other alternatives or by changing the wording. A basic conclusion of these investigators was that linguistic confusion was not the source of the conjunction fallacy in the Linda problem.

The conjunction fallacy is of particular interest to us since it appears to be a logical error rather than a shortcut or an error due to a lack of information. Some of the other biases and heuristics follow quite naturally from normal learning and memory processes. According to the "availability" heuristic, for instance, people estimate the probability of events by the ease of recalling information about those events. The probabilities of sensational events are overestimated since they receive considerable press coverage. The availability heuristic, then, is consistent with standard theories about learning and cognition. In contrast, most people have had considerable experience with uncertain events and their eventual outcomes, and one might expect them to learn, at least implicitly, that joint events are less probable than the constituent events.

We carried out four experiments to examine aspects of the conjunction fallacy. In the first experiment, we attempted to make the scenario more realistic and concrete for the subjects. On the basis of the outcome of the first experiment, we began to question whether the correct rank

of the compound alternative might depend on what the subject assumes about the context in which the alternatives are embedded; in particular, is the outcome of the scenario known and might it be included as one of the alternatives? In the remaining experiments, we used scenarios in which the context was specified more explicitly.

EXPERIMENT 1

In a number of problem-solving contexts, subjects fail to solve a problem when it is stated in an abstract form but do much better when the problem is made more concrete. For example, the majority of subjects fail to solve an "if/then" problem when stated abstractly (Wason & Johnson-Laird, 1972) but readily solve a conceptually similar problem when it is phrased in terms of letters and postage stamps (Johnson-Laird & Wason, 1978) or in terms of drinking age and ordering alcoholic beverages (Griggs & Cox, 1982). We wondered if concreteness would affect the conjunction fallacy in a similar fashion. The Linda problem is relatively concrete, and Tversky and Kahneman argued that subjects do better on an abstract version of the problem, although they provide no data to support this assertion. We wondered, however, whether an even more concrete and familiar scenario might alter students' judgments.

Method

Materials. We used three scenarios in the first experiment: the Linda problem described above, a scenario concerning a student's course schedule for the quarter, and a scenario about a horse race. We felt that all students would be highly familiar with constructing and thinking about course schedules and most of the students would be familiar with horse races. Both of those scenarios provided natural ways of describing compound events. The latter two scenarios are listed below.

Class Schedule

On the way to the Registrar's Office a few days before winter term, the registrar accidentally dropped three students' schedule cards into the slush. After retrieving the cards, the registrar saw that the ink had smeared, making the cards almost totally unreadable. When the cards had dried, the registrar was able to make out one or two of the courses on each card. When you are in the Registrar's Office, you happen to overhear that Ben Parker, a junior, was one of the students whose card was involved in the mishap. Because he is a pretty good friend of yours, you volunteer to see if you can figure out which one is Ben's.

Ben had been off last term on an internship in Washington DC. In fact, he has had several internships relating to his future career goals. Two of his other internships were in Washington DC as an aide for a New York senator. These internships gave him a great deal of exposure to the Senate proceedings. For one of his leave terms, he was offered an internship in the Finance Office of the White House. He would have loved to work in the White House, but because of his lack of interest in anything mathematical, he turned the offer down. This winter, Ben will be back at Dartmouth taking classes.

The registrar shows you the three smeared schedule cards and asks you to number the cards in order from most likely (1) to least likely (3) to be Ben's schedule.

— Schedule 1:	Government 34	Math 23	?
— Schedule 2:	?	?	Government 34
— Schedule 3:	Math 23	?	?

Horse Race

The third race at Churchill Downs is a stakes race featuring nine 3-year-olds. The race is over a mile and an eighth and the track is in good condition. We are going to focus on two of the horses entered in the race. The first one, *HardRock*, has run seven previous races and only finished in the money once, and that was over a muddy track. According to the handicapper for the Daily Racing Form, *HardRock* is picked to finish seventh. *Avenger* is the favorite for the third race. He has finished in the money seven times in nine races, winning five of them. He is well rested and appears to be in excellent form for today's race. Listed below are three betting tickets for the third race; each ticket costs \$2 and pays \$6 if successful. Please rank order the tickets with (1) being the one that you would most like to have and (3) being the one that you would least like to have.

- Ticket 1: *HardRock* to finish second.
- Ticket 2: *Avenger* to win.
- Ticket 3: *Avenger* to win and *HardRock* to finish second.

Subjects and Procedure. These scenarios were distributed to the students in an introductory psychology class containing 130 students, one scenario per student. The two courses mentioned in the schedule scenario are taught at Dartmouth College and the subjects were familiar with the approximate content of the courses. All possible orders for the three alternatives were equally represented for each scenario. The students were asked to read the scenario and rank the alternatives.

Results and Discussion

The results of the first experiment appear in Table 1. The percentage of students making the conjunction fallacy on the Linda problem (77%) was about the same as that found by Tversky and Kahneman (1983). Contrary to our expectations, the two new scenarios led to opposite outcomes. The students who received the schedule scenario were more likely to make the conjunction fallacy than were the students judging the Linda problem. The difference between the two conditions was significant [$\chi^2(1) = 8.42, p < .01$]. The horse race problem led to significantly fewer conjunction errors than either of the other conditions.

We were surprised at the outcome in the schedule condition. We had thought that it was the most concrete of the three scenarios and the one with which the students would be most familiar. On reflection, however, we realized that ranking the compound event ahead of the unlikely single event might not have been a fallacy in the schedule problem. This is because we know that Ben's schedule does exist and it is likely to be one of the ones that we are shown. In this situation, a Bayesian analysis becomes appropriate and the different alternatives represent different forms of evidence. In the horse race problem, the event has not yet occurred and we have no

idea whether any of the alternatives listed will be correct. A more detailed analysis of this distinction follows.

Models for Unknown versus Known Outcomes

We are suggesting that different models are appropriate for analyzing the horse race and schedule problems. The outcome of the horse race is not yet known and it might be quite different from any of the alternatives listed. Ben's schedule is stated to be one of the listed alternatives and the subjects' task is to figure out which one.

Model for unknown outcomes. This model applies to events in which we are trying to predict an unknown outcome. It is the model that applies to most gambling situations. With this model, the probability of a joint event is always equal to or less than that of the component events. If the two component events are independent, then the probability of the joint event will be the product of the probabilities of the two events. If the two component events are not independent, then the probability of the joint event can range from zero (disjoint events) to the probability of the component event with the smallest probability, depending on the nature of the dependence between the two component events. When this model is appropriate, it is definitely a fallacy to rate the likelihood of a conjunction higher than the likelihood of either constituent event.

Model for known outcomes. In some situations, though, we know that an outcome has occurred and our job is to discover the correct outcome. For instance, a robbery might have occurred in a small town and the local police chief knows all of the possible suspects. She first learns that dynamite was used in the robbery. Person *X* is unlikely to use dynamite, so the police chief does not place *X* very high on the suspect list. The police chief then discovers that the robber had bright red hair. Since *X* has bright red hair, he moves to a much higher spot on the suspect list. This is an appropriate way to revise the probability of *X*'s having committed the crime and is consistent with the use of Bayes's theorem.

The key to ranking suspects in this problem is not the relative probability of the joint and constituent events. It is the probability of each of those events for *X* relative to the probability of those events for the other suspects. The probability of the joint event "red hair and used dynamite," given *X*, is still less than or equal to the probability of the "used dynamite," given *X* as with the model for unknown events, but the joint event might well be more characteristic of *X* relative to the other suspects than one of the constituent events is of *X* relative to the other suspects. In other words, we are dealing with conjunctive evidence rather than conjunctive events.

We can evaluate the likelihood of each of the alternatives in the schedule problem by using Bayes's theorem and choosing some plausible probabilities for each of the necessary events. The scenario was formulated so that Ben was quite likely to take Government 34 but unlikely to

Table 1
Percentage of Subjects Who Made the Conjunction Fallacy in Experiment 1

Condition	Percentage Making Fallacy	<i>N</i>
Linda-Control	77	44
Schedule	98	44
Horse Race	52	42

Note— $\chi^2(2) = 24.38, p < .01$.

take Math 23. Some possible parameter values are listed below:

$$P(\text{Gov 34} | \text{Ben}) = .20.$$

$$P(\text{Math 23} | \text{B}) = .05.$$

$P(\text{G34 and M23} | \text{B}) = .02$ (do not assume independence).

$P(\text{B}) = .33$ (this is the prior since we are told that Ben is one of three).

$$P(\text{G34} | \text{not B}) = .05.^1$$

$$P(\text{M23} | \text{not B}) = .15.$$

$$P(\text{G34 and M23} | \text{not B}) = .01.$$

$$P(\text{not B}) = .67.$$

With these parameters, we can use Bayes's theorem to estimate the probabilities of the two critical alternatives:

$$\begin{aligned} P(\text{B} | \text{M23}) &= \frac{P(\text{M23} | \text{B})P(\text{B})}{P(\text{M23} | \text{B})P(\text{B}) + P(\text{M23} | \text{not B})P(\text{not B})} \\ &= \frac{(.05)(.33)}{(.05)(.33) + (.15)(.67)} = .14 \end{aligned}$$

and

$$\begin{aligned} P(\text{B} | \text{M23} \& \text{G34}) &= \frac{P(\text{M23} \& \text{G34} | \text{B})P(\text{B})}{P(\text{M23} \& \text{G34} | \text{B})P(\text{B}) + P(\text{M23} \& \text{G34} | \text{not B})P(\text{not B})} \\ &= \frac{(.02)(.33)}{(.02)(.33) + (.01)(.67)} = .50 \end{aligned}$$

So, for the particular values chosen, the compound alternative is more probable than one of the component alternatives.² How important are the values of the various parameters? The critical parameters are the relative probabilities of the various events for Ben versus "not Ben." As long as the ratio of the compound event for Ben to the compound event for "not Ben" is larger than the ratio of one of the component events for Ben versus "not Ben," then the compound alternative will have a higher likelihood than will that component event. In fact, any ordering of the three alternatives is possible, depending on the set of parameter values. For the typical problem with one likely event and one unlikely event, the model will often predict the compound alternative to be more probable than the unlikely event alone.

The prior probability of Ben is not relevant, as long as it is defined. This is important because it means that we do not have to be told or infer that one of the alternatives is necessarily the correct one. All that is necessary is that we know or infer that the correct alternative is known and that each of the listed alternatives has some a priori probability of being that known outcome. In the horse race problem, the outcome has not yet occurred,

so the prior probabilities of the various alternatives are undefined.

What about the extensional property? This property does not apply in the schedule scenario, since each alternative is considered a unique schedule that has more or less information on it. The alternative specifying a single event does not include the alternative specifying the joint event; it represents someone else's schedule.

What about the Linda problem described earlier? The Linda scenario seems neutral with respect to the models outlined above. Subjects are not told that Linda's fate is known or that one of the alternatives represents Linda. Subjects are not told the contrary, either. It seems to us that the most appropriate framing would be to assume that the outcome is unknown and that the unknown model applies, but subjects easily could assume that the three alternatives represent three different individuals and that one of the alternatives represents Linda. In the latter case, rating the joint event higher than bank teller alone does not represent a fallacy.

There are data in Tversky and Kahneman's (1983) study that provide some support for a distinction between the two models or ways of viewing the problem. They used a variety of different scenarios, some of which were ambiguous with respect to the appropriate model, and some of which were clearly described as unknown future events or gambling situations. The latter unambiguously required the unknown model. The ambiguous problems included the Linda, Bill, and "dyspnea" scenarios. Across various groups of subjects, the percentage of subjects committing the conjunction fallacy for the Linda scenario ranged from 85% to 90%; for the Bill scenario, it ranged from 83% to 92%; and for the dyspnea scenario, 91% of their subjects made a conjunction error. Two scenarios in which the unknown probability model clearly applies are the "Bjorn Borg" tennis scenario and the "RGRRR colored die" scenario. The percentage of subjects making the conjunction error for the Borg problem was 72%, and for the colored die problem, it ranged from 62% to 65%. The differences in the percentages of conjunction errors for the ambiguous problems (Linda, Bill, and dyspnea) versus the unambiguous problems (Borg and colored die) is highly significant, given the numbers of subjects used. There were other scenarios in Tversky and Kahneman's study, but most of them explicitly manipulated some additional variable, such as causality, making the scenarios difficult to classify.

There is an additional finding in Tversky and Kahneman's (1983) study contrasting these two types of scenarios that is pertinent to the distinction that we are making. They (Tversky & Kahneman, p. 299) presented subjects directly with a choice between two arguments that described the principles involved in the Linda problem, and asked which argument was the most convincing. Only 35% of the subjects chose the argument describing the conjunction (or extensional) rule. They presented a different group of subjects with a similar choice concerning the

colored die problem (p. 304). With the colored die problem, 76% of the subjects chose the extensional argument. The difference between the two percentages is highly significant. We would argue that in the Linda problem, the subjects might have interpreted the scenario in a fashion in which the extensional argument does not apply, but that the colored die problem does not afford such an interpretation. In the remaining experiments, we attempted to manipulate subjects' interpretations of the Linda problem by embellishing the scenario.

EXPERIMENT 2

In the second experiment, we used three versions of the Linda problem: the original version described above (henceforth, the "control version"), a version that emphasized that Linda's fate was known and that one of the three alternatives represented her, and a version that made it clear that Linda's fate would not be known for several years and that the goal was to predict it. We expected the explicitness of the known/unknown distinction to influence the subjects' likelihood ratings. In particular, we expected more conjunction errors with the known scenario and fewer conjunction fallacies for the unknown scenario, relative to the control scenario.

Method

Materials. The two new versions of the Linda problem appear below:

Linda Known

Blake Electronics needed to hire a new public relations manager. They had managed to narrow their list of applicants down to the three most qualified. Now the decision is in the hands of the Personnel Manager, Barry Green. His staff had taken care of narrowing the pool of applicants down to three, and now it was his job to interview the candidates and make a decision. He thought it would be a good idea to look over the folders of the three applicants before the interviews began. As he sat down at his desk, he accidentally spilled a cup of coffee. When the spilled coffee was finally wiped up, Barry noticed that most of the information on the applications had been smeared and could no longer be read. Barry decided to ask his staff to try to remember as much about the applicants as they could. You are called from your department in Blake Electronics because the staff recalled that you knew one of the applicants. Although you hadn't seen her for some time, Linda was one of your good friends from college days. You agree to try and figure out which of the applications is Linda's based on the information that can still be read.

Linda is now 31. You remember her as single, outspoken, and very bright. She majored in philosophy in college. As a student, she was deeply concerned with discrimination and other social issues, and participated in anti-nuclear demonstrations.

The staff asks you to rank order the following applications with (1) being the most likely and (3) being the least likely. (---- indicates information that was covered by coffee.)

	Occupation	Social interests	Hobbies
___ Application 1:	Bank teller	----	----
___ Application 2:	----	Active in the feminist movement	----
___ Application 3:	Bank teller	Active in the feminist movement	----

Linda Unknown

Small-town colleges often have rituals for the graduating seniors to go through before they pass out into the ranks of the "real world." One small college had a tradition of predicting what the seniors would

be like in 10 years. The fun of this tradition is to read the predictions at the 10-year class reunion. The predictions are generated at random with a computer program written at the college. The computer sometimes predicts the occupation that the student will have 10 years from now, sometimes predicts an avocation or cause that the person will be involved with, and sometimes both. Some predictions turn out to be fairly accurate, but many times the predictions are off the wall. You are shown three predictions generated for your roommate, Linda, and asked to place a bet on the one that you think is most likely to come true.

In 10 years Linda will be 31. You know her as single, outspoken, and very bright. She majored in philosophy. She was deeply concerned with discrimination and other social issues, and participated in anti-nuclear demonstrations.

Below are the three predictions generated for Linda. Please rank order the predictions from (1) being the most likely to (3) being least likely to be true of Linda 10 years from now.

- ___ Prediction 1: Linda will be a bank teller.
- ___ Prediction 2: Linda will be active in the feminist movement
- ___ Prediction 3: Linda will be a bank teller and active in the feminist movement.

Subjects and Procedure. The subjects for this experiment were drawn from various psychology courses at Dartmouth. All of the courses were beyond the introductory level and most of the students had had at least one other psychology course. There were 135 subjects, 45 for each of the three scenarios. All possible orders for the three alternatives were represented for each of the scenarios.

Results and Discussion

The data for Experiment 2 appear in Table 2. The results were in the predicted direction and the differences were highly significant [$\chi^2(2) = 13.35, p = .0014$].

Even though the results were in the predicted direction, the majority of subjects in the unknown group made the conjunction fallacy, and in this condition it is definitely a fallacy. We wondered, though, whether statistical training might make a difference when the underlying conditions are clearly specified. As mentioned earlier, Tversky and Kahneman (1983) reported little or no effect of such training on the frequency of making the conjunction fallacy with the direct technique used here. (They did find an effect of training in what they referred to as the "transparent" version.) Most of the studies on training, however, use the Linda or the Bill problem, and we argue that those problems are ambiguous with respect to the two models outlined earlier. If the experts are interpreting those scenarios as representing known outcomes, then their behavior is appropriate.

Unfortunately, we did not collect any information concerning probability knowledge from our subjects until late in the experiment and have relatively complete information for only the unknown scenario. We asked those subjects after they had rated the alternatives to indicate whether they had had a course in, or significant exposure to, probability theory. Twenty-three of the subjects in the unknown condition answered "yes" and 22 answered "no." The data, broken down by experience, appear in Table 2. The subjects with some exposure to probability theory made significantly fewer conjunction fallacies than did those without such exposure [$\chi^2(1) = 3.94, p = .04$]. Because of our sketchy data concerning probability background, we carried out an additional experiment to examine the effect of training.

Table 2
Percentage of Subjects Who Made the Conjunction
Fallacy in Experiments 2 and 3

Educational Experience	Condition		
	Linda-Control	Linda-Known	Linda-Unknown
Undergraduates*			
Naive			73% (22)
Some Statistics			43% (23)
Combined	82% (45)	89% (45)	58% (45)
Medical Students†			
1st Year	92% (12)	100% (12)	65% (17)
2nd Year	89% (19)	100% (13)	41% (17)
Combined	90% (31)	100% (25)	53% (34)

Note—Numbers in parentheses refer to number of subjects in that cell. $*\chi^2(2) = 13.35, p < .01$. $\dagger\chi^2(2) = 22.87, p < .01$.

EXPERIMENT 3

Medical students at the Dartmouth Medical School receive a segment on decision theory as part of their second-year studies. The segment occurs in the middle of the year. We had the 41 first- and 49 second-year medical students participate in the present experiment subsequent to the occurrence of that year's decision-theory segment. The second-year students, but not the first-year students, therefore had had instruction on decision theory as medical students. Each medical student received one of the three scenarios used in Experiment 2.

The overall results of the experiment and the results as a function of training appear in Table 2. The overall results were similar to those obtained in Experiment 2. The effect of scenario was highly significant [$\chi^2(2) = 22.87, p \approx 0$]. The effect of training was not significant for any of the scenarios, but the results for the unknown scenario mirrored those obtained in Experiment 2.

If we look at the students in Experiments 2 and 3 who had had some exposure to probability or decision theory, we see for the first time that the majority of those students did *not* make the conjunction fallacy in the one condition in which it is clearly a fallacy. Lehman, Lempert, and Nisbett (1988) showed that exposure to graduate training in psychology and medicine led to improved reasoning on problems from a similar domain.

EXPERIMENT 4

Tversky and Kahneman (1983) reported that phrasing a question as a bet reduced the likelihood of making a conjunction fallacy. We believed that it might not be betting, per se, that reduced the number of conjunction errors. Instead, it seemed possible that the mention of betting encouraged the subjects to adopt a future/unknown interpretation and that it was this future/unknown interpretation that led to fewer conjunction fallacies.

Each of our previous problems with a future/unknown interpretation had inadvertently included betting as part of the scenario. We wanted to know whether it was the future/unknown interpretation or the betting that led to

the reduced percentage of conjunction errors. We predicted that there would be no significant difference between a future/unknown scenario that included betting and one that did not include betting.

Method

Materials. In Experiment 4, we used two versions of the Linda Unknown problem. One version made obvious the fact that the person would be betting money. In the second version, all mention of betting was removed from the scenario.

Subjects. We used 48 subjects—24 for each version of the scenario—from an introductory psychology class at Dartmouth.

Results

The results of the experiment appear in Table 3. The percentage of conjunction fallacies was not significantly different from previous experiments using the Linda Unknown problem. There was no significant difference between the betting and the no-betting scenarios. Admittedly, we are accepting the null hypothesis in this experiment, but the results are actually in the opposite direction from those obtained by Tversky and Kahneman (1983) using a different scenario.

GENERAL DISCUSSION

We showed that, in certain contexts, an alternative describing the conjunction of two events can be more probable than an alternative describing one of the constituent events. In particular, this is true if the outcome of a scenario is or can be assumed to be known and might be one of the alternatives listed. In such a case, each alternative is judged relative to the likelihood of that event in the population at large, and it is more appropriate to talk about conjunctive evidence than the extensional rule in those cases. We argue that many of the scenarios used in previous research did not specify this aspect of the problem clearly. If subjects assume that the outcome is known, then they may not be making a fallacy.

We asked a group of 12 psychology graduate students and faculty members to rate the control versions of the Linda and Bill problems as to whether the outcomes were known or unknown. We used a 7-point scale in which 1 indicated that the outcome was already known and might be one of the alternatives listed and 7 indicated that the outcome would not be known until sometime in the future. The Linda problem received a mean rating of 2.4 and the Bill problem a rating of 2.6. These ratings indicated that, for most of our subjects, the known model would be appropriate, and rating the conjunction higher

Table 3
Percentage of Subjects Who Made the Conjunction
Fallacy in Experiment 4

Condition	Percentage Making Fallacy	<i>N</i>
Linda Unknown-Bet	66.7	24
Linda Unknown-No Bet	62.5	24

Note— $\chi^2(1) = 0.0911, p = .76$.

than one of the constituents would not necessarily be a fallacy.

We described data from Tversky and Kahneman (1983) and from our own research suggesting that subjects are sensitive (at least implicitly) to the distinction between the two models and the types of situations to which they apply.

It should be emphasized that even though we have shown some variables that influence the likelihood of a conjunction error, many subjects in the most favorable conditions continue to make the conjunction fallacy. This may imply, as Tversky and Kahneman suggested, that many people misunderstand basic probability theory and are influenced by representativeness. On the other hand, it may indicate that many people have a strong tendency to invoke the "known model," even when that model is inappropriate.³ If the latter possibility is true, it might be due to our considerable experience with situations in which the outcomes are known and our task is to search for clues as to the correct alternative. Further research is required to understand the remaining, and frequent, instances of the conjunction fallacy.

Our data suggest that any training program on decision theory should consider the distinction between the two types of models and describe the situations in which they are appropriate.

REFERENCES

- EDWARDS, W. (1968). Conservatism in human information processing. In B. Kleinmuntz (Ed.), *Formal representation in human judgment* (pp. 17-52). New York: Wiley.
- GRIGGS, R. A., & COX, J. R. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, *73*, 407-420.
- JOHNSON-LAIRD, P. N., & WASON, P. C. (Eds.) (1978). *Thinking*. Cambridge: Cambridge University Press.
- KAHNEMAN, D., & TVERSKY, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237-251.
- LEHMAN, D. R., LEMPERT, R. O., & NISBETT, R. E. (1988). The effects of graduate training on reasoning. *American Psychologist*, *43*, 431-442.
- 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.
- LYON, D., & SLOVIC, P. (1976). Dominance of accuracy information and neglect of base rates in probability estimation. *Acta Psychologica*, *40*, 287-298.
- MORIER, D. M., & BORGIDA, E. (1984). The conjunction fallacy: A task specific phenomenon? *Personality & Social Psychology Bulletin*, *10*, 243-252.
- TVERSKY, A., & KAHNEMAN, D. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 84-98). New York: Cambridge University Press.
- TVERSKY, A., & KAHNEMAN, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, *90*, 293-315.
- WASON, P. C., & JOHNSON-LAIRD, P. C. (1972). *Psychology of reasoning*. Cambridge, MA: Harvard University Press.

NOTES

1. In these equations, "not Ben" refers to the probability of a randomly chosen Dartmouth student other than Ben taking one of the indicated courses.

2. The Bayesian model is an oversimplification for the schedule scenario as described. We presented the model because it is essentially correct and easier to appreciate. The model would be appropriate if the subjects were asked to consider a particular alternative in isolation. The subjects, however, were shown all three alternatives at the same time. In this situation, it is necessary to consider a sample space consisting of six 3-tuples representing all of the possible ways of assigning the three alternatives to Ben and two unknown, randomly chosen individuals. The priors for each 3-tuple would be identical. It is possible to compare the predictions from the two models. For the Bayesian model, it is necessary to compute the probability of each alternative using the parameters listed and then to normalize those probabilities to sum to one. With this approach, the probability of choosing the alternative containing Government 34 alone is .55, for Math 23 alone is .09, and for the compound event is .36. Using the more complete model, the respective probabilities are .51, .11, and .38. Both models, as described, refer to only the first alternative chosen, but either model could easily be extended to predict the entire rank ordering. Both models, however, predict the compound event to be more likely than the unlikely single event, and the predictions of the two models are quite close.

3. We are basically arguing here that people intuitively prefer a Bayesian approach to these problems. This may seem to be at odds with considerable research showing that people are "bad Bayesians." In many situations, they tend to ignore base rates in estimating probabilities (Kahneman & Tversky, 1973; Lyon & Slovic, 1976), and in other situations, they are too conservative in revising their estimates on the basis of new information (Edwards, 1968). For our purposes, however, subjects do not have to use Bayes's theorem well as long as they base their likelihoods on the ratios of the probabilities for the events in question to the probability of the same events for the population at large.

(Manuscript received July 7, 1988;
revision accepted for publication May 18, 1989.)