NOTICE

The quality of this microform is heavily dependent upon the quality of the original thesis submitted for microfilming. Every effort has been made to ensure the highest quality of reproduction possible.

If pages are missing, contact the university which granted the degree.

Some pages may have indistinct print especially if the original pages were typed with a poor typewriter ribbon or if the university sent us an inferior photocopy.

Previously copyrighted materials (journal articles, published tests, etc.) are not filmed.

AVIS

La qualité de cette microforme dépend grandement de la qualité de la thèse soumise au microfilmage. Nous avons tout fait pour assurer une qualité supérieure de reproduction.

S'il manque des pages, veuillez communiquer avec l'université qui a conféré le grade.

La qualité d'impression de certaines pages peut laisser à désirer, surtout si les pages originales ont été dactylographiées à l'aide d'un ruban usé ou si l'université nous a fait parvenir une photocopie de qualité inférieure.

Les documents qui font déjà l'objet d'un droit d'auteur (articles de revue, tests publiés, etc.) ne sont pas microfilmés.

THE UNIVERSITY OF ALBERTA

THE SUBJECTIVE PROBABILITY OF COMPOUND EVENTS: WHY DO PEOPLE COMMIT THE CONJUNCTION FALLACY?

BY

IGOR GAVANSKI

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES AND RESEARCH

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF PSYCHOLOGY

EDMONTON, ALBERTA

FALL, 1988

ISBN 0-315-45452-0

THE UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR: IGOR GAVANSKI

TITLE OF THESIS: THE SUBJECTIVE PROBABILITY OF COMPOUND EVENTS: WHY DO PEOPLE COMMIT THE CONJUNCTION FALLACY?

DEGREE: DOCTOR OF PHILOSOPHY

YEAR THIS DEGREE GRANTED: FALL, 1988

(SIGNED) ....

PERMANENT ADDRESS:

833 Woodbridge Drive

Bloomington, Indiana

U.S.A. 47408

Date: 26 September 1988

THE UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled THE SUBJECTIVE PROBABILITY OF COMPOUND EVENTS: WHY DO PEOPLE COMMIT THE CONJUNCTION FALLACY? submitted by Igor Gavanski in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

..........Gary L. Wells..........
Supervisor

..........Craig A. Anderson..........

..........Curt Hoffman..........

..........Kellogg Wilson..........

..........[signature]..........

Date: ..26 September...... 1988

Abstract

People commonly violate a basic rule of probability theory, judging a compound event to be more probable than at least one of its component events. This error is called the conjunction fallacy, and has been attributed to people's reliance on the representativeness heuristic for judging probability (e.g., Tversky & Kahneman, 1983). This research evaluated two alternative interpretations of the conjunction fallacy. First, it was proposed that people's natural representations of information and habitual reasoning strategies result in systematic preferences for certain sample spaces, and that the conjunction fallacy occurs when problems evoke inappropriate sample spaces. Second, it was proposed that the conjunction fallacy occurs because people judge conjunctions by separately assessing the probability of each component event and then combining these probabilities using incorrect combination-rules. The first two experiments examined the inappropriate-sample-space interpretation. Experiment 1 tested the hypothesis that problems requiring judgments of the probability that an instance belongs to a conjunction of categories lead to conjunction fallacies because people preferentially use categories as sample spaces. Although manipulating the instance-category structure of problems had the predicted effects on the sampling strategies subjects used, there was no evidence that the fallacy results from inappropriate sample spaces. Experiment 2 tested the hypothesis that conjunctions of dependent events yield conjunction fallacies because people treat such conjunctions as conditionals. This hypothesis was supported, with the results demonstrating that both positive dependencies [i.e., $P(B/A) > P(B)$] and negative dependencies [i.e., $P(B/A) < P(B)$] can lead to conjunction fallacies. Experiments 3 and 4 examined the incorrect-combination-rule interpretation of the fallacy. These experiments provided strong evidence that the fallacy is largely the result of incorrect combination-rules, and is related to representativeness only insofar as representativeness influences the component-event probabilities.

ACKNOWLEDGEMENTS

I am very grateful to Gary Wells for his guidance and encouragement throughout my doctoral program and as dissertation supervisor. I also wish to thank the members of my dissertation committee, Craig Anderson, Curt Hoffman, Richard Johnson, and Kellogg Wilson for their many helpful comments.

TABLE OF CONTENTS

LIST OF TABLES

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 anti-nuclear demonstrations. What is the probability of each of the following?

A. Linda is active in the feminist movement.
B. Linda is a bank teller.
C. Linda is a bank teller and is active in the feminist movement.

People presented this problem commonly commit the conjunction fallacy, judging the compound event C to be more probable than B. This judgment is readily recognized as a fallacy when one considers that the set of bank tellers includes both bank tellers who are active in the feminist movement and those who are not. The conjunction fallacy is a reliable phenomenon in problems of this kind, having been replicated with various problems and subject populations (e.g., Crandall & Greenfield, 1986; Morier & Borgida, 1984; Tversky & Kahneman, 1982; 1983; Wells, 1985).

Tversky and Kahneman (e.g., 1982, 1983) have provided the most comprehensive and widely accepted explanation of the conjunction fallacy, attributing it to people's reliance on the representativeness heuristic for judging probability. According to this explanation, subjective probabilities are based on assessments of representativeness, or similarity. Similarity is described as a feature-matching process, in which the similarity between objects is a linear combination of their common and distinctive features (Tversky, 1977). Thus, when judging the probability that object B belongs to class A, or that process A will generate event B, people assess how similar B is in its essential features to A. When B is very similar to A, then its probability is judged to be high. Because similarity does not obey the laws of probability, reliance on the representativeness heuristic sometimes leads to systematic errors such as the conjunction fallacy. In the Linda problem, for example, Linda's description is more similar to feminist bank tellers than to bank tellers, so people judge the conjunction, *bank teller and feminist*, as more likely than its component, *bank teller*.

One source of evidence that the conjunction fallacy is caused by representativeness-mediated judgments is research showing that the occurrence of the fallacy depends on the representativeness relations among target descriptions and the component events of conjunctions. For example, Tversky & Kahneman (1983) found that the conjunction fallacy rarely occurred in the Linda problem when subjects were told only that Linda is 31 years old. Wells (1985) systematically varied the representativeness of the component events of conjunctions and found that the fallacy occurred most frequently when a representative and unrepresentative event were paired (average = 72%), and least frequently with pairings of two unrepresentative events (average = 7%).

Tversky and Kahneman (1980) also found a very close correspondence between people's judgments of probability and of representativeness. They gave subjects conjunction problems, each containing a target description followed by a list of occupational and political categories or conjunctions of these. They found a correlation of .96 between mean rankings of the categories according to how representative the target description was of them and mean rankings of the categories according to the probability that the target belongs to them.

Although these data are consistent with the argument that the conjunction fallacy occurs because people judge probability according to representativeness, they are also consistent with other interpretations. Most importantly, it is unclear exactly what representativeness means to subjects and what is the direction of the causal relation between representativeness and probability. Although it is undoubtedly the case that A is similar to B if A possesses many features characteristic of B's stereotype, it is also the case that A will be similar to B if many A's are B's or many B's are A's. In other words, assessments of similarity may sometimes be based on judgments of relative frequency or probability, i.e., P(B/A), P(A/B), rather than forming the basis for these judgments.

Inappropriate Sample Spaces and the Conjunction Fallacy

In this section, I will develop the argument that the conjunction fallacy occurs, not because of people's reliance on representativeness, but because the task requirements conflict with people's natural representations of the information provided in the problems. To correctly solve conjunction problems, people must form appropriate sample spaces. I suggest that in many problems that yield a high incidence of conjunction fallacies, the appropriate sample space is difficult to create and the problem structure tends to naturally evoke an inappropriate sample space.

The idea that inappropriate sample spaces may contribute to conjunction fallacies is not new. Some theorists have argued that semantic ambiguities in the problems may lead subjects to form inappropriate sample spaces. Markus and Zajonc (1985) suggested that "many subjects may have simply misunderstood the [conjunction] problem and formed the wrong sample space *because* of their misunderstanding" (p. 190). They proposed that subjects may misinterpret the component events A and B of a conjunction as [A & not-B] and [B & not-A]. In the Linda problem, for example, subjects may interpret the statement "Linda is a bank teller" as "Linda is a bank teller who is not active in the feminist movement". If so, then subjects' judgments of conjunctive probabilities do not necessarily violate probability theory; although P(A & B) cannot be higher than P(A) or P(B), it may well be higher than P(A & not-B) or P(B & not-A). Morier and Borgida (1984) found that structuring problems to avoid this kind of semantic confusion reduced (but did not eliminate) the incidence of conjunction fallacies for one of two problems they presented subjects. Other research addressing the issue of semantic ambiguity has ruled out many, although not all, possible sources of semantic confusion as primary causes of the conjunction fallacy (see Tversky & Kahneman, 1983; Crandall & Greenfield, 1986).

The present conceptualization, however, does not rely on semantic ambiguity to explain why subjects might form inappropriate sample spaces. Rather, I propose that people's natural representations of information and their habitual reasoning strategies result in systematic preferences for certain sample spaces. Inappropriate sample spaces may therefore be used even in the absence of semantic ambiguity.

I am not suggesting that people solve conjunction problems by systematically defining a sample space and performing extensive mental calculations to determine the relative frequency of instances in the sample space matching a particular profile. Rather, as with most cognitive tasks, people may use relatively simple heuristic strategies such as sampling instances or exemplars according to availability (cf., Tversky & Kahneman, 1973), and then assessing the match between the exemplars and the profile according to representativeness. I do suggest, however, that it may not be reliance on these heuristic strategies per se that causes conjunction fallacies, but rather, that these occur because people sample from inappropriate sample spaces.

Different problems may evoke different sample spaces. The following sections describe three general classes of conjunction problems that seem particularly likely to evoke inappropriate sample spaces. Although I do not present this classification as exhaustive, most problems yielding a high incidence of the conjunction fallacy appear to fall within one or more of these classes.

Problems evoking instance-category representations. Many conjunction problems used in previous research share a common structure. Subjects read a description of a target person and judge the probability that the target is a member of certain categories or classes [e.g., P(Bank teller/Description), P(Feminist/Description)] or of conjunctions of classes [e.g., P(Bank Teller & Feminist/Description)]. The correct sample space in problems of this sort is *people of the target's description*. Such problems, however, seem particularly likely to evoke an incorrect sample space. Because the target is described in highly specific terms (e.g., name, age, personal or educational history), he or she is naturally represented as a unique instance. So it may be difficult to form a sample space of people fitting the target's description. Categories, on the other hand, readily serve as sample spaces. Categories are used to partition instances, so that nonidentical instances can be treated as equivalent (Wattenmaker, Dewey, Murphy & Medin, 1986). Any category (e.g., bank tellers, feminists) contains numerous instances, so it is relatively easy to form a sample space comprising instances within the category. Furthermore, because instances are assigned to categories rather than vice versa, it may be simpler to assess what proportion of category members fit a target description than what proportion of instances fitting the target description are members of a category. In the former case, all relevant instances are contained within one category, whereas in the latter case, relevant instances may be spread across several categories and thus be more difficult to retrieve. For these reasons, people may be particularly prone to use categories as sample spaces in problems evoking instance-category representations, and to sample instances from the categories to determine how well they match the target description -- in effect judging P(Description/Category) rather than P(Category/Description). In the Linda conjunction problem, for example, people may sample instances or exemplars of bank tellers and feminists to determine what proportion of them match Linda's description [P(Description/Bank teller & Feminist)], rather than sample instances from people of Linda's description to determine what proportion are bank tellers and feminists [P(Bank teller & Feminist/Description)]. This strategy would account for the occurrence of the conjunction fallacy, because there is no reason that P(Description/Bank teller & Feminist) must be lower than P(Description/Bankteller), and indeed it may reasonably be considered to be higher.

Consistent with the inappropriate-sample-space interpretation, Tversky and Kahneman (1983) found that the incidence of the conjunction fallacy was substantially reduced when subjects were provided a sample of people fitting the target description and judged the relative frequency of targets for which the conjunction would be true. Providing a sample of people fitting the target description may make it easier for subjects to form the appropriate sample space by reducing the likelihood that the target will be represented as a unique instance, and having subjects make relative frequency judgments clearly directs them to the appropriate sample space. However, it is unclear at this point whether this procedure is effective because it causes subjects to shift their sample space or simply because it causes them to forego reliance on representativeness in favour of an extensional strategy (cf. Tversky & Kahneman, 1983).

<u>Problems conflicting with natural direction of inference.</u> Occupational or other consistent demands often require people to make inferences or predictions primarily in one direction. For example,

physicians generally diagnose illness based on presenting symptoms, i.e., P(Illness/Symptoms). When problems require people to make judgments that conflict with their natural or habitual direction of inference [e.g., requiring people to judge P(B/A) when they habitually judge P(A/B)], they may be particularly prone to use an incorrect sample space (e.g., a sample space of instances of B, rather than instances of A). Consider the following problem: Mr. Jones has Disease X. Symptom A is highly representative of the disease and symptom B is not representative of the disease. What is the probability that Mr. Jones has: Symptom A? Symptom B? Symptom A and Symptom B? Tversky and Kahneman (1983) presented problems of this general structure to practicing physicians and found that the vast majority (73 to 100%) committed the conjunction fallacy. To correctly solve these problems, physicians would have to form a sample space of *people who have the disease* and assess what proportion of these would have the symptoms. However, because inferences from diseases to symptoms conflict with practitioners' usual direction of inference, they may tend to consider the symptoms as *given*, and sample *people who have the symptoms* to determine what proportion have the disease. In other words, physicians may be assessing P(Disease/Symptoms) rather than assessing P(Symptoms/Disease) (c.f., Eddy, 1982; Einhorn & Hogarth, 1986). This strategy would account for the conjunction fallacy in problems like the above, because P(Disease/Representive Symptom & Unrepresentative Symptom) is clearly higher than P(Disease/Unrepresentative Symptom).

Problems evoking conditional interpretations. People may sometimes interpret a conjunctive probability as a conditional probability. Tversky and Kahneman (1983) have suggested that conjunctions in which one component is a cause of the other (e.g., John smoked & John died of cancer) yield conjunction fallacies because "it is more natural to assess the probability of the effect given the cause than the joint probability of the effect and the cause" (p. 308). Cause and effect relations entail a "positive dependency" between events (Falk & Bar-Hillel, 1983), because the occurrence of one event (the cause) increases the probability of the other (the effect). It may not be only problems involving positive dependencies that evoke conditional interpretations. Conditional interpretations of conjunctions may also be evoked by negative dependencies between the components of the conjunction (i.e., when the occurrence of one event reduces the probability of the other). Such interpretations seem particularly likely when the two components have a clear temporal order, in which case people may interpret P(First event & Second event) as P(Second event/First event) (cf., Braine, 1979). Indeed, some conjunction fallacies that have been attributed to reliance on representativeness may well be due to conditional interpretations of this sort. Consider the following problem that Tversky and Kahneman presented to subjects shortly after Bjorn Borg had won his fifth Wimbledon tournament:

> Suppose Bjorn Borg reaches the Wimbledon finals in 1981. Please rank order the following outcomes from most to least likely.
>
> A. Borg will win the match
>
> B. Borg will lose the first set
>
> C. Borg will lose the first set but win the match
>
> D. Borg will win the first set but lose the match

Of subjects presented this problem, 72% committed the conjunction fallacy, assigning a higher rank to C than to B. Tversky and Kahneman attributed this to the fact that C was more representative of Borg than was B. This problem, however, meets the criteria outlined above that may evoke conditional interpretations: the two components of the conjunction have an explicit temporal order and there is a strong (negative) dependency between them. Thus, subjects may solve this problem by forming a sample space of *matches in which Borg has lost the first set*, and judging the probability that he will then go on to win the match [i.e., P(Win match/ Lost first set)]. Although this probability will fall below P(Win match) it may well be higher than P(Lose first set). Tversky and Kahneman reject this interpretation because subjects who were presented with sequences of Borg's wins and losses (e.g., LWWLW) and asked which were consistent with the statement, "Borg will lose the first set but win the match," endorsed sequences consistent with the conjunctive interpretation and not other interpretations of the statement. However, it is unclear how this procedure would distinguish between conjunctive and conditional interpretations of the statement. Furthermore, as argued earlier, people may form inappropriate sample spaces even when they do not consciously misinterpret the problem. The inappropriate sample space interpretation therefore remains a viable alternative to the representativeness interpretation of problems of this sort.

Incorrect Combination-Rules and the Conjunction Fallacy

Some conjunction fallacies do not readily lend themselves to either the sample-space or the representativeness explanation. For example, people sometimes commit conjunction fallacies when simply given various combinations of the prior and conditional probabilities of the conjunction's component events and they base their estimates of the conjunctive probability on this information (e.g., Beyth-Marom, 1981; Fischoff, 1985; Goldsmith, 1978; Slovic, 1969; Wyer, 1970; 1976). The conjunction fallacy in such problems does not depend on the specific nature of the component events, or on representativeness-mediated probability judgments, because it occurs even when symbolic events are used (Beyth-Marom, 1981). People cannot judge a conjunction of symbolic events directly, but must assess its probability by somehow combining the prior and/or conditional probabilities of its component events. Thus, conjunction fallacies in these problems must be due to the use of incorrect rules for combining the probabilities of the conjunction's component events.

Could the conjunction fallacy in representativeness problems also be due to these incorrect combination-rules? Recall that the conjunction fallacy is found primarily when a representative anc unrepresentative event are paired (Wells, 1985). Because representativeness and subjective probability are highly correlated (Tversky & Kahneman, 1980; 1983), this pairing results in one high- and one low-probability event. Many incorrect combination-rules suggested in the literature (see Goldsmith, 1978; Wyer, 1976; Wyer & Goldberg, 1970; Abelson, Leddo, & Gross, 1987) are most likely to result in conjunction fallacies, or in the largest magnitude of the fallacy, when the probabilities of the two component events are substantially different (as when a high- and a low-probability event are paired). Indeed, data presented by Beyth-Marom (1981; Table 2) yield a correlation of .57 between the mean difference in estimated probabilities of the conjunction's two components [i.e., P(A) - P(B), where B is the less probable component] and the mean magnitude of the conjunction fallacy [i.e., P(A & B) - P(B)].

Thus, the relation between representativeness and the conjunction fallacy may be at least partly artifactual: Representativeness may be related to the conjunction fallacy only to the extent that it affects the subjective probabilities of component events, which people then combine using incorrect rules. These incorrect rules tend to yield conjunction fallacies with particular combinations of component event probabilities, regardless of whether these probabilities were derived through representativeness or other means (e.g., the probabilities were simply given in a conjunction problem).

According to this incorrect combination-rule interpretation, people do not judge the probability of the conjunction directly. Rather, they decompose the conjunction into its component events, judge the probabilities of the components according to representativeness, and then combine these probabilities to arrive at the probability of the conjunction. But why would people judge the probabilities of the component events directly but use a different strategy in judging the probability of the conjunction? The answer is that, particularly when a conjunction is composed of a representative and an unrepresentative event, there is generally no natural category or stereotype corresponding to the conjunction that would allow for a representativeness judgment. Consider the Linda problem. People certainly have stereotypes for feminists and stereotypes for bank tellers against which Linda's description can be compared, but do they have stereotypes for feminist bank tellers? This seems highly unlikely, as the two categories share few common attributes and are somewhat incompatible. Therefore, subjects cannot assess Linda's similarity to a stereotypic category of feminist bank tellers, but must separately assess similarity to feminists and to bank tellers, resulting in two similarity and two probability judgments, which must somehow be combined. In most research using conjunctions of a representative and an unrepresentative event, the two components of the conjunction are at least somewhat incompatible, or unrepresentative of each other (e.g., bank teller & feminist; accountant & plays jazz for a hobby; lose first set & win match). In general, it may be that whenever the components of a conjunction are unrepresentative of each other or differ in their similarity to the target description, people tend to decompose the conjunction and make separate judgments for each component.

Research Objectives

The goals of this research were to determine why people commit the conjunction fallacy, and in particular to evaluate the adequacy of the representativeness interpretation of the fallacy, relative to the inappropriate-sample-space and the incorrect-combination-rule interpretations. Experiments 1 and 2 were designed to examine the role of inappropriate sample spaces in the conjunction fallacy. Experiment 1 examined whether people preferentially use categories as sample spaces in conjunction problems, and whether inappropriate sample spaces can account for the conjunction fallacy in problems evoking instance-category representations. Experiment 2 examined whether conditional interpretations of conjunctions involving negative or positive dependencies between the two components contribute to the fallacy. Experiments 3 and 4 were designed to determine to what extent the fallacy occurs because people use incorrect rules for combining the probabilities of component events to arrive at the probability of the conjunction.

Experiment 1

Inappropriate Sample Spaces and the Conjunction Fallacy

I have proposed two reasons that conjunction problems evoking instance-category representations may lead people to form inappropriate sample spaces: First, people may have difficulty forming the appropriate sample space because they view the target as a unique instance; second, people may preferentially use categories as sample spaces because of the way they naturally represent information about categories and instances.

Based on this conceptual framework, conjunction problems were constructed to evoke either inappropriate or appropriate sample spaces. Each problem contained a description of a target person; the conjunction was composed of one event that was representative of the description and one that was unrepresentative of the description. The problems varied according to: (a) whether the target was described as a specific instance or as a category member and whether the conjunction was composed of categories or of instances of behavior, and (b) whether or not the target was drawn from a pool of people fitting a common description. According to the inappropriate-sample-space explanation, the conjunction fallacy should occur most frequently when the target is described as an instance and the conjunction is composed of categories, least frequently when the target is described as a category member and the conjunction is composed of instances of behavior, and with intermediate frequency when neither the target description nor the conjunction strongly evoke category representations (because this problem structure should not bias subjects toward using either appropriate or inappropriate sample spaces). Providing a sample of people sharing the target description should reduce the incidence of the conjunction fallacy by making subjects less likely to perceive the target as a unique instance, and thus making it easier for them to form the appropriate sample space.

Method

Subjects and Design

Subjects were 123 male and female introductory psychology students who participated in partial fulfillment of a course requirement. Each subject was randomly assigned to one condition of a 2 X 3 (Sample Provided vs. No-sample Provided X Problem Structure) between-subjects factorial design. From 18 to 23 subjects participated in each condition.

Stimulus Materials

There were three kinds of problem structures: instance-category, instance-instance, and category-instance. Four problems of each structure were constructed or adapted from previous research (see Appendix A). Each conjunction was composed of one event that was representative and one that was unrepresentative of a target description.

Instance-category problems. These problems correspond structurally to those used in much previous research, in which subjects were given a specific description of a target person's background and behavior. The two components of the conjunction were social or occupational categories and subjects judged the probability that the target instance belongs to these categories (hence the designation instance-category). Following is an example of an instance-category problem:

> Jane is in her late 20's and divorced. She is assertive and outspoken. She is university-educated, well dressed and works for a large successful corporation. Her career requires that she travel frequently. What is the probability of each of the following?
>
> A. Jane is an executive
>
> B. Jane is a mother
>
> C. Jane is both an executive and a mother

Instance-instance problems. These contained the same target descriptions as the instance-category problems, but the conjunctions were composed of instances of behavior (or behavioral intentions and preferences) rather than of natural categories; subjects judged the probability that the target instance would engage in these instances of behavior. Following is an example of these problems:

> Jason is in his late 20's, single and lives alone. He has travelled extensively and always enjoys seeing new things. He likes to go to clubs to hear new bands; he writes poetry and plays the guitar. What is the probability of each of the following?
>
> A. Jason has unusual taste in clothing.
>
> B. Jason wants to marry a traditional woman.
>
> C. Jason both has unusual taste in clothing and wants to marry a traditional woman.

Category-instance problems. In these problems, the target was simply identified as belonging to a social or occupational category (e.g., feminist, bank teller, musician, nun); as with the instance-instance problems, the conjunction was composed of instances of behavior and subjects judged the probability that the category member would engage in these instances of behavior. Following is an example of these problems:

> Marie is a Roman Catholic nun. What is the probability of each of the following?
>
> A. Marie enjoys spending time alone with her thoughts.
>
> B. Marie spends much of her free time watching T.V.
>
> C. Marie both enjoys spending time alone with her thoughts and spends much of her free time watching T.V.

Sample vs. no-sample manipulation. There were two versions of each problem. The no-sample versions were as described above. The sample versions were identical except that they contained information that the target was randomly selected from a large pool of people fitting a common description. The target descriptions for these problems were modified to read as follows (bracketed portions varied between problems):

For purposes of a psychological study, researchers have selected a large pool of men (women) from across the country who share many common characteristics. All of these men (women) are [Target Description]. One person, [Target Name], has been randomly selected from this group.

Strategy Questionnaires

A strategy questionnaire was constructed for each problem. The questionnaire required subjects to indicate on an 8-point scale (0 = did not use this strategy at all; 7 = used this strategy very heavily) how heavily they had relied on each of four strategies in judging the probability of each component event and of the conjunction. Appendix B contains a sample strategy questionnaire, which was used for the Linda instance-category problem. Strategies 1.1, 1.6, and 1.11 were the appropriate sampling strategies for judging P(A), P(B), and P(A & B), respectively. Strategies 1.2, 1.7, and 1.12 were inappropriate sampling strategies. In judging the probability that *Linda is a bank teller*, for example, the appropriate strategy is to sample *people of Linda's description* to determine what proportion are bank tellers (1.1). An inappropriate strategy is to sample *bank tellers* to determine what proportion match Linda's description (1.2). However, what appears to be an inappropriate sampling strategy may in fact be part of a normatively appropriate Bayesian analysis. According to Bayes' Theorem,

$$P(\text{Bank teller/Description}) = \frac{P(\text{Description/Bank teller})P(\text{Bank teller})}{P(\text{Description})}$$

Thus, if subjects who report using the "inappropriate" sampling strategy [i.e., P(Description/Bank teller)] are using anything like a Bayesian analysis, then they would also have to estimate the prior probability of bank teller. Strategies 1.3, 1.8, and 1.13, are prior probability strategies, included to determine if the inappropriate sampling strategy could be part of a Bayesian analysis. Strategies 1.4, 1.9, and 1.14 are representativeness strategies, included to assess subjects' reliance on representativeness in judging probabilities.

Procedure

Subjects participated in groups of 4 to 16. Each subject received a question booklet containing instructions and four conjunction problems with corresponding strategy questionnaires, and an answer booklet. There were two problem orders, each the reverse of the other. Subjects were told that the purpose of the study was to examine people's intuitions about probability. They were instructed to answer the problems in the order that they appeared, to answer all the questions on each page before looking at the next page, and not to change their anwers once they had recorded them. Subjects made their probability estimates as percentages ranging from 0% to 100%, with 0% meaning that "the statement is definitely untrue," 50% meaning that "the statement is equally likely to be true or untrue," and 100% meaning that "the statement is definitely true". For each problem, subjects first judged the probability of the conjunction's component events and then the probability of the conjunction. On the page following each problem subjects filled out a strategy questionnaire indicating how heavily they had relied on each of several strategies in making their probability judgments.

After subjects had judged all four problems, they estimated the base-rate probabilities of the conjunction's two component events A and B for each problem (i.e., the probability of the events in the general population). These probability estimates were requested as follows: "What is the probability that a man [woman] drawn at random from the general population would A [B]? Comparisons of these prior probability estimates to subjects' estimates of the conditional probabilities of the events, given the target description, provide a manipulation check to determine if the representative and unrepresentative events really were perceived as such by subjects. Subjects then judged the conditional probability of A/B or B/A for each problem. Half the subjects in each condition judged P(A/B) and half judged P(B/A). These probability estimates were requested as follows: "If you already know that a man [woman] drawn at random from the general population A [B], what then is the probability that he [she] also B [A]?

The last page of the booklet contained a background questionnaire designed to assess subjects' training in, and understanding of, probability theory (see Appendix C). The questionnaire included questions about the correct combination rules for judging the probability of conjunctions composed of independent events and disjunctions composed of mutually exclusive events, about the definition of mutually exclusive events, and about how much formal education the subject had in probability theory or statistics. The last question required subjects to make an intuitive assessment of their overall knowledge of probability theory on an 8-point scale anchored with "very poor" and "excellent".

Results

Table 1 displays the mean subjective probabilities of conjunctions and representative (R) and unrepresentative (U) component events for subjects in all conditions. As can be seen in this table, the subjective probability of the representative event was substantially higher than that of the unrepresentative event for each problem. The mean estimates of P(R) were 71.4, 65.0, and 71.9 for instance-category, instance-instance, and category-instance problems, respectively; the mean estimates of P(U) were 27.8, 25,9, and 18.5, respectively. However, an event may be judged probable or improbable simply because it has a high or low base-rate frequency in the general population, and not because of its representativeness relation with a target description. To more directly determine the effectiveness of the representativeness manipulation, I compared subjects' estimates of the base-rate probabilities of the events with their estimates of the probabilities of the events given the target descriptions. The base-rate probability should be lower than the target probability for representative events, and higher for unrepresentative events. This was the case. For representative events, the mean difference between the target probability and the base-rate probability was 43.2, 32.5, and 32.3 for instance-category, instance-instance, and category-instance problems, respectively. For unrepresentative events, the mean difference was -7.0, -7.0, and -8.5 for the three problem structures, respectively. Thus, subjects did perceive the representative and unrepresentative events as

Table 1

Mean Subjective Probabilities of the Conjunction (R&U) and its Representative (R) and Unrepresentative (U) Events For Experiment 1

	Problem 1			Problem 2			Problem 3			Problem 4			Overall		
	R	U	R&U	R	U	R&U	R	U	R&U	R	U	R&U	R	U	R&U
Instance-Category															
No Sample	79.5	31.0	46.9	79.0	22.7	29.6	84.5	25.2	33.4	67.9	19.0	19.5	77.7	24.5	32.4
Sample	69.0	25.6	27.0	58.6	30.0	24.8	76.7	41.4	41.6	57.8	23.2	15.6	65.5	30.1	27.3
Instance-Instance															
No Sample	52.9	24.8	23.5	62.1	26.7	32.9	79.2	22.8	40.5	63.7	28.9	38.8	64.5	25.8	33.9
Sample	42.2	24.6	19.2	65.0	40.8	40.8	81.1	18.3	24.9	73.3	20.6	34.4	65.4	26.1	29.8
Category-Instance															
No Sample	81.9	18.3	30.3	70.2	14.3	23.6	76.0	19.3	21.2	62.4	22.3	32.2	72.7	18.6	26.8
Sample	73.3	12.8	21.8	74.4	13.1	19.7	77.6	22.5	28.7	53.1	25.3	18.6	69.6	18.4	22.2

such, and the representativeness of these events was similar for the three problem structures.

Because conditional dependencies between the component events of a conjunction may affect the incidence of the fallacy (see Experiment 2), it is important to determine if the dependencies differed substantially across problem structures. This was not the case. The conditional probabilities of each event given the other were lower than the base-rate probabilities of the events for all problem structures. The mean difference between the base-rate probabilities and the conditional probabilities was 9.9, 9.4 and 11.5 for instance-category, instance-instance, and category-instance problems, respectively. Thus, there were slight negative dependencies between component events for all three problem structures, and the dependencies were similar across problem structures.

Strategy questionnaire results. Table 2 displays the mean ratings of each strategy according to how heavily it was used for the probability judgments. Mean ratings of each strategy were analysed in 3 X 2 (Problem Structure X Sample vs. No Sample) analyses of variance. The problem structure manipulation had the predicted effects on sampling strategies. As seen in Table 2, The inappropriate sampling strategy was used most heavily by instance-category subjects (for whom the category was the inappropriate sample space) and least heavily by category-instance subjects (for whom the category was the appropriate sample space), $F(2, 117) = 3.26$, $p < .05$. The opposite pattern emerged for the appropriate sampling strategy, which was used most heavily by category-instance subjects and least heavily by instance-category subjects, $F(2, 117) = 7.89$, $p < .001$. Instance-instance subjects made intermediate use of both the appropriate and inappropriate sampling strategies. The mean difference between ratings for the appropriate and inappropriate sampling strategies (i.e., appropriate strategy - inappropriate strategy) significantly differed for all three problem structures, $ts > 2.19$, $ps < .05$. The sample vs. no-sample manipulation had no significant effect on subjects' sampling strategies, either as a main effect, $ps > .50$, or in interaction with problem structure, $ps > .10$.

Table 2

Mean Ratings of Strategy Use in Experiment 1

	Problem Structure			
	Instance-Category	Instance-Instance	Category-Instance	Overall
Inappropriate Sampling Strategy	3.3	2.9	2.4	2.9
Appropriate Sampling Strategy	4.0	4.3	5.2	4.5
Prior Probability Strategy	1.7	2.4	2.6	2.2
Representativeness Strategy	3.5	4.0	4.4	4.0

Note. Strategy ratings for judgments of the component events and the conjunction were averaged for each problem and then across problems.

The inappropriate sampling strategy does not appear to have been part of a Bayesian analysis. If subjects who relied on this strategy were using a Bayesian analysis, then they should also have relied heavily on the prior-probability strategy. In fact, however, the prior probability strategy was used least by instance-category subjects (who relied most on the inappropriate sampling strategy) and most by category-instance subjects (who relied least on the inappropriate sampling strategy), $F(2, 117) = 5.24$, $p < .01$. Furthermore, within each problem structure, subjects who relied most on the inappropriate sampling strategy (by a median split) relied least on the prior-probability strategy, $ts > 2.09$, $ps < .05$.

Subjects' reliance on representativeness was influenced by both the problem structure manipulation, $F(2, 117) = 3.28$, $p < .05$, and the sample manipulation, $F(2, 117) = 4.23$, $p < .05$. As can be seen in Table 2, the representativeness strategy was used most heavily by category-instance subjects and least heavily by instance-category subjects. Subjects provided with a sample fitting the target description relied less on representativeness than subjects not provided a sample.

As would be expected, reliance on the representativeness strategy was positively correlated to the mean subjective probabilities of representative component events, $r(121) = .34$, $p < .001$ and negatively correlated to the mean probabilities of unrepresentative events, $r(121) = -.17$, $p < .05$. [2]

Interestingly, subjects in all conditions relied on the two sampling strategies and the representativeness strategy less for judging the conjunctive probability than for judging the component-event probabilities. The difference between mean strategy ratings for component events and for conjunctions was significant greater than zero for all three of these strategies, paired difference $ts (122) > 7.89$, $ps < .001$. This pattern occurred for each problem structure and attained significance for all three strategies within each problem structure, $ts > 2.76$, $ps < .01$.

Incidence and Magnitude of the Conjunction Fallacy

Two measures of the conjunction fallacy were calculated for each subject and separate analyses were performed on each measure. The incidence of the fallacy was simply the proportion of problems resulting in conjunction fallacies for each subject. Thus, this index varied from .00 to 1.00 in increments of .25. The magnitude of the fallacy was the difference between the conjunctive probability and the lower of the two component probabilities for each problem (cf. Wells, 1985). Positive differences mean the conjunctive probability exceeded the lower component probability; negative differences mean the lower component probability exceeded the conjunctive probability. Difference scores were calculated for each of the four problems that the subject judged; these were then averaged to form one magnitude-of-fallacy measure for each subject.

Table 3 displays the mean magnitude and incidence of the conjunction fallacy for subjects in all conditions. Problem structure had no significant effect on either the magnitude of the fallacy, $F(2, 117) = 1.09$, $p > .30$, or on the incidence of the fallacy, $F(2, 117) = 1.58$, $p > .20$. Providing subjects a sample of people fitting the target description did significantly reduce the magnitude of the fallacy, $F(1, 117) = 9.05$, $p < .01$, and resulted in a marginally significant reduction in the incidence of the fallacy, $F(1, 117) = 3.28$, $p = .07$. However, as can be seen in Table 3, this manipulation had much stronger effects on both the magnitude and incidence of the fallacy for instance-category problems than for instance-instance or category-instance problems. A contrast analysis showed that this interaction pattern was significant for both the magnitude and incidence measures, $Fs(1, 117) > 8.60$, $ps < .01$, and that the residual variance was not significant, $Fs < 1$.

However, the sample manipulation also influenced the subjective probabilities of component events for instance-category problems, but not for instance-instance or category-instance problems. As can be seen in Table 1, providing a sample decreased P(R) and increased P(U) for instance-category problems, but had little effect on these for the other two problem structures. Analysis of variance on the mean difference between P(R) and P(U) yielded a significant Problem Structure X Sample interaction, $F(2, 117) = 3.12$, $p < .05$; individual contrasts showed that the sample effect was significant for instance-category problems, $t(117) = 3.26$, $p < .001$, but not for the other two problem structures, $ts(117) < 1$. Because the difference between P(R) and P(U) was significantly

Table 3

Mean Magnitude and Incidence of the Conjunction Fallacy for Experiment 1

	Problem 1	Problem 2	Problem 3	Problem 4	Overall
Instance-Category					
No Sample	15.9 (.57)	7.0 (.57)	8.2 (.43)	.5 (.24)	7.9 (.45)
Sample	1.5 (.29)	-5.3 (.19)	.2 (.24)	-7.7 (.10)	-2.8 (.20)
Instance-Instance					
No Sample	-1.3 (.32)	6.2 (.42)	17.7 (.58)	9.8 (.47)	8.1 (.45)
Sample	-5.4 (.39)	0.0 (.39)	6.6 (.33)	13.9 (.56)	3.8 (.42)
Category-Instance					
No Sample	12.0 (.65)	9.3 (.48)	1.8 (.30)	9.9 (.48)	8.3 (.48)
Sample	9.0 (.52)	6.6 (.38)	6.2 (.43)	-6.7 (.38)	3.8 (.43)

Note. The magnitude of the conjunction fallacy is the conjunctive probability minus the lower of the two component probabilities. Thus, positive numbers mean that the conjunctive probability was higher than at least one component probability; negative numbers mean that the conjunctive probability was lower than both component probabilities. Number in parentheses is the incidence of the conjunction fallacy.

correlated with the magnitude and incidence of the conjunction fallacy, rs(123) = .24 and .44, respectively, ps < .01, the effect of the sample manipulation on the fallacy for instance-category problems may have been due to its effects on the component-event probabilities. Analyses of covariance, controlling for the difference between P(R) and P(U), supported this interpretation. With the difference between P(R) and P(U) as a covariate, the sample manipulation had no significant effects on either the incidence or magnitude of the fallacy for instance-category problems, ps > .05. Including the covariate reduced the proportion of variance accounted for by the sample manipulation (r^2) from about .20 to less than .07 for both measures of the fallacy.

Were subjects' strategy reports related to the incidence of the conjunction fallacy? This was the case only for the representativeness strategy. Reliance on this strategy was positively correlated with both the magnitude and incidence of the fallacy, magnitude r(121) = .22; incidence r(121) = .28, ps < .001. However, partial correlations, controlling for the difference between P(R) and P(U), did not attain significance, partial rs(120) < .14, ps > .10. Thus, reliance on representativeness and the incidence of the fallacy may be related primarily because of their common relation with the component-event probabilities.

Did judging several problems, or filling out a strategy questionnaire following each problem, make subjects either more or less likely to commit the conjunction fallacy? This was examined by comparing the incidence of the fallacy for the two problem orders. Half the subjects in each condition judged Problem 1 first (before they had seen any strategy questionnaires or other problems) and Problem 4 last, and half judged the problems in the reverse order. There were no significant order effects either for problems that served as Problem 1 or for those that served as Problem 4, ts < 1.[3] The mean incidence of conjunction fallacies for Problems 1 and 4 was .40 when they were presented first and .43 when they were presented last. Thus, the incidence of the fallacy was unaffected by whether or not subjects had previously judged several problems and filled out strategy questionnaires.

Responses to the background questionnaire showed that only a minority of subjects knew the correct combination-rule for judging conjunctive probabilities (27%) or disjunctive probabilities (20%). However, correct answers to the background questions (sum of questions 1 to 3) were only weakly related to the incidence of the fallacy, $r(119) = -.22$, $p < .05$. The same was the case for subjects' level of training in probability theory and statistics, $r = -.13$, $p < .20$, and their self-ratings of their knowledge of probability theory, $r = -.17$, $p < .10$.

Discussion

Manipulating the instance-category structure of problems had the predicted effects on subjects' reported sampling strategies. Subjects were most likely to use the appropriate sampling strategy [i.e., P(A & B/target description)] and least likely to use the inappropriate strategy [i.e., P(target description/A & B)] when the target was described as a category member and the conjunction was composed of instances of behavior. This pattern was reversed when the target was described as an instance and the conjunction was composed of categories. When neither the target description nor the conjunction strongly evoked category representations, subjects' reliance on both the appropriate and inappropriate sampling strategies was intermediate. Although mean ratings of the appropriate strategy were higher than those of the inappropriate strategy in all conditions, it is likely that subjects' ratings provide a conservative estimate of their reliance on the inappropriate strategy. Because the two strategies were juxtaposed on the questionnaires, some subjects may well have recognized which was appropriate, and adjusted their ratings accordingly.

There was clear evidence that the inappropriate sampling strategy was not part of a Bayesian analysis, because subjects who relied most on this strategy relied least on the prior probability strategy. This pattern not only rules out a Bayesian analysis, but also further validates subjects' self-reports, since the prior probabilities are relevant to the appropriate, but not the inappropriate, sampling strategy. For example, the prior probability of bank teller is relevant to a judgment of P(bank teller/target description). But if subjects instead judge P(target description/bank teller), then bank teller is given and its prior probability is irrelevant.

The strategy questionnaire results provide good evidence that people do have systematic preferences for certain sample spaces, and in particular, that they tend to form sample spaces from categories. There was, however, no evidence that such sampling preferences contributed to the conjunction fallacy. Subjects in all conditions committed the fallacy about equally, with the exception of the instance-category subjects who were provided with a sample fitting the target description. Analyses of covariance showed that the low incidence of the fallacy for these subjects was largely attributable to the effect of the sample manipulation on the subjective probabilities of the component events. Providing a sample sharing the target description reduced the subjective probability of the representative event and increased that of the unrepresentative event for instance-category subjects, but not for the other two groups. This effect in turn may have been due to the effect of the sample manipulation on subjects' reliance on representativeness. Reliance on representativeness was positively correlated to P(R) and negatively correlated to P(U), and the instance-category subjects who were provided with a sample relied less on representativeness than subjects in other conditions. It is unclear why the sample manipulation had this effect for instance-category subjects, because it did not significantly influence their use of other strategies.

Although the inappropriate-sample-space interpretation of the conjunction fallacy was not supported, these findings provide only equivocal support for the representativeness interpretation of the fallacy. Reliance on representativeness was related to the fallacy, but there was no evidence that representativeness influenced the fallacy independently of its effects on the component-event probabilities. Furthermore, subjects relied less on representativeness (and on both sampling strategies) for judging conjunctive probabilities than for judging component probabilities, suggesting that they may have used some other strategy to judge conjunctions. The implications of this finding will be further explored in the general discussion.

[illegible]

well be more probable [illegible]

the conditional interpretation [illegible]

[illegible]

match or [illegible] the [illegible]

In the present study [illegible]

conjunction problem that was [illegible]

In the present study, subjects were each given one version of a conjunction problem that was conceptually similar to the Bjorn Borg problem. The conjunction was composed of two high-probability outcomes (HH), two low-probability outcomes (LL), or one high- and one low-probability outcome (HL or LH, depending on which outcome occurred first). The outcomes were either dependent or independent: the dependency was positive for HH and LL pairings, and negative for HL and LH pairings. If people interpret the conjunction as a conditional when the component events are dependent but not otherwise, then changing the temporal order of high- and low-probability events (i.e., HL vs. LH) should affect the incidence of the conjunction fallacy only for dependent-event conjunctions. Specifically, LH pairings should yield more fallacies than HL pairings for dependent-event conjunctions but LH and HL pairings should yield equal numbers of fallacies for independent-event conjunctions. Furthermore, dependent-event conjunctions should yield more fallacies than independent-event conjunctions for HH, LL, and LH pairings, but fewer fallacies for HL pairings.

Method

Subjects and Design

The sample was composed of 131 male and female introductory students who participated in partial fulfillment of a course requirement. Subjects were randomly assigned to conditions of a 2 X 2 X 2 (First Outcome X Second Outcome X Outcome Dependency: dependent vs. independent) between-subjects factorial design. From 15 to 18 subjects participated in each condition.

Stimuli

Because the Bjorn Borg problem is outdated (i.e., Borg is no longer among the top-ranked players in the world), eight versions of a conceptually similar conjunction problem were created. The problem involved a playoff series between two National Hockey League teams, the Edmonton Oilers and the Vancouver Canucks. The Oilers consistently finish the regular season at or near the top of the league standings and, at the time this research was conducted, they had won the Stanley

Cup finals in four of the last five years. The Vancouver Canucks, on the other hand, consistently finish the season in the bottom half of the league standings, and had not defeated the Oilers in 24 consecutive games between the two teams. Thus, it is highly representative (and probable) for the Oilers to win against the Canucks and highly unrepresentative (and improbable) for them to lose. The problem read as follows:

> Suppose the Edmonton Oilers and the Vancouver Canucks meet in the Smythe Division Final next season. The Smythe Division Final is a best-of-seven series. What is the probability of each of the following outcomes?

In independent-outcome problems, conjunctions were composed of the outcomes of the first and second games of the series (e.g., The Oilers will win the first game and the Oilers will lose the second game). Although these two outcomes may not be completely independent, it was expected that most subjects would not consider them to be strongly dependent. Four independent-event problems included all possible permutations of outcomes of the first two games (i.e., win-win, win-lose, lose-win, lose-lose). In dependent-outcome problems, conjunctions were composed of the outcome of the first game and the final outcome of the series (e,g., The Oilers will win the first game and the Oilers will lose the series). Four dependent-event problems included all possible permutations of outcomes of the first game and the series.

Procedure

Subjects participated in small groups. Introductory instructions were as in Experiment 1. Subjects each received one version of the Oiler-Canuck problem. They first judged the probability of each outcome and then the probability of the conjunction of outcomes. On the following page, subjects judged the conditional probability, P(Second Outcome/First Outcome). These judgments were elicited by questions of the following form: "Suppose you already know that the

Oilers have won [lost] the first game, what then is the probability that they will go on to win [lose] the second game, [series]."

Results

Manipulation checks. Table 4 displays mean subjective probabilities of the conjunction and its component outcomes, as well as the conditional probability of the second outcome given the first, for subjects in all conditions. As can be seen in Table 4, the manipulation of outcome-probability was successful. Subjects considered it substantially more likely for the Oilers to win than to lose against the Canucks, regardless of whether the outcome was of a game or of the entire series. To evaluate the success of the manipulation of outcome dependency, I compared subjects' estimates of the probability of the second outcome, $P(O_2)$, with their estimates of the conditional probability of the second outcome, given the first outcome, $P(O_2/O_1)$. The two probabilities should be identical if the outcomes are seen as independent, $P(O_2)$ should be higher than $P(O_2/O_1)$ if the outcomes are seen as negatively dependent, and lower if the outcomes are seen as positively dependent. In independent-outcome conditions (i.e., first game and second game), 65% of subjects' estimates of $P(O_2)$ were identical to their estimates of $P(O_2/O_1)$, whereas in dependent-outcome conditions (i.e., first game & series), only 42% of these probabilities were identical, $^2 = 6.48$, $p < .02$. As can be seen in Table 4, the directions of the dependencies were as expected, with win-win and lose-lose being positive, and win-lose and lose-win being negative.

The conjunction fallacy. The incidence of the conjunction fallacy was coded by assigning subjects a 1 if they committed the fallacy and a 0 otherwise. Although this measure is binomially distributed, the normal distribution provides a good approximation of the binomial with sample sizes such as in this study (Hays, 1981). Furthermore, even with highly skewed distributions, the sampling distribution of the mean rapidly approaches normality with increasing sample size. The incidence of the fallacy was therefore analysed using standard

Table 4

Mean Subjective Probabilities of the Conjunction ($O_1 \& O_2$) and the Prior and Conditional Probabilities of its Component Events (O_1, O_2, and O_2/O_1) for Experiment 2

	Component Outcomes of the Conjunction															
	Win-Win				Win-Lose				Lose-Win				Lose-Lose			
	O_1	O_2	O_2/O_1	$O_1 \& O_2$	O_1	O_2	O_2/O_1	$O_1 \& O_2$	O_1	O_2	O_2/O_1	$O_1 \& O_2$	O_1	O_2	O_2/O_1	$O_1 \& O_2$
Independent (Game 1 & Game 2)	73.1	70.9	70.3	59.4	68.7	[illegible]	27.5	37.1	25.5	76.4	75.4	34.5	27.0	31.3	31.7	22.3
Dependent (Game 1 & Series)	69.7	74.0	76.5	67.3	79.2	14.3	11.1	14.6	34.6	82.2	68.6	50.9	30.6	19.8	21.9	16.6

Note. O_1 and O_2 are the temporally first and second outcomes, respectiveless, that comprised the conjunction.

parametric tests. The magnitude of the fallacy was calculated as in Experiment 1. Table 5 displays the mean magnitude and incidence of the fallacy for subjects in all conditions. Consistent with Wells' (1985) finding, the conjunction fallacy occurred more frequently with pairings of a high- and a low-probability event (win-lose; lose-win) than with pairings of high-probability events (win-win) or of low-probability events (lose-lose). Analyses of variance on the incidence and magnitude of the fallacy yielded significant effects for outcome-combination, incidence $F(3,123) = 10.22$, $p < .001$, magnitude $F(3,123) = 7.72$, $p < .001$. A contrast comparing the win-lose and lose-win conjunctions with the win-win and lose-lose conjunctions showed that both the incidence and magnitude of the fallacy were significantly higher for the former conjunctions, incidence $t(1,123) = 4.65$, $p < .001$; magnitude $t(1, 123) = 4.18$, $p < .001$.[6]

Separate a priori contrasts were applied to test the predicted interaction between outcome combination and outcome dependency for negative- and for positive-dependency conditions. Negative dependencies were hypothesized to decrease the incidence of the fallacy for win-lose conjunctions, but increase it for lose-win conjunctions. Table 5 shows that the results conform to this pattern: dependent-outcome problems resulted in fewer fallacies than independent-outcome problems for win-lose conjunctions (.27 vs. .47), but more fallacies for lose-win conjunctions (.78 vs. .47). The incidence of the fallacy did not differ between win-lose and lose-win conjunctions for independent-outcome conditions (.47 vs. .47). The following contrast weights compared the incidence and the magnitude of the fallacy for the two negative dependency cells with the corresponding independent-outcome cells: for independent outcomes, -1 (win-lose) -1 (lose-win); for dependent outcomes, -2 (win-lose) 4 (lose-win). This contrast attained significance for both the incidence and magnitude measures, incidence $t(123) = 3.45$, $p < .001$; magnitude $t(123) = 2.37$, $p < .01$.

Positive dependencies were hypothesized to increase the incidence of the conjunction fallacy. As can be seen in Table 5, the two positive dependency conditions (i.e., win game & win series; lose game

Table 5

Magnitude and Incidence of the Conjunction Fallacy in Experiment 2

	Component Outcomes of the Conjunction			
	Win-Win	Win-Lose	Lose-Win	Lose-Lose
Independent (Game 1 and Game 2)	-10.1 (.00)	9.9 (.47)	9.1 (.47)	-3.3 (.13)
Dependent (Game 1 and Series)	-1.0 (.20)	.3 (.27)	16.3 (.78)	-.2 (.28)

Note. The magnitude of the conjunction fallacy is the conjunctive probability minus the lower of the two component probabilities. Thus, positive numbers mean that the conjunctive probability was higher than at least one component probability; negative numbers mean that the conjunctive probability was lower than both component probabilities. Number in parentheses is the incidence of the conjunction fallacy.

& lose series) led to a higher incidence of the conjunction fallacy than the corresponding independent-outcome conditions (.24 vs. .06). A contrast comparing the positive dependency cells with the corresponding independent-outcome cells was significant for the incidence of the fallacy, $t(123) = 1.67$, $p < .05$, but did not attain significance for the magnitude of the fallacy, $t(123) = 1.45$, $p = .08$.

If the observed effects of outcome-dependency occurred because dependent outcomes caused subjects to treat the conjunction as a conditional, then estimates of the conjunctive probability should be more similar to estimates of the conditional probability for dependent-outcome subjects than for independent-outcome subjects. To test this, I calculated the absolute difference between $P(O_1 \& O_2)$ and $P(O_2/O_1)$ for all subjects. The mean of these difference scores was significantly lower for dependent-outcome subjects (mean = 10.89) than for independent-outcome subjects (mean = 21.22), $F(1, 123) = 9.89$, $p < .01$. Furthermore, the conjunctive and conditional probabilities were identical for a higher proportion of dependent-outcome subjects than independent-outcome subjects (45% vs. 23%), $^2 = 7.27$, $p < .01$.

Discussion

The manipulation of outcome dependency had the predicted effects on the conjunction fallacy. Positive dependencies increased the incidence of the fallacy. Negative dependencies either increased or decreased the incidence of the fallacy depending on the temporal order of the high- and low-probability outcomes. Conjunctions of a low probability outcome followed by a high probability outcome (i.e., lose-win) yielded a lower incidence of the fallacy for dependent than for independent outcomes, whereas conjunctions of a high probability outcome followed by a low probability outcome (i.e., win-lose) yielded a higher incidence of the fallacy for dependent outcomes. The temporal order of high- and low-probability outcomes had no effect on the incidence of the fallacy for independent-outcome conditions.

Consistent with the argument that these effects occurred because subjects tended to treat conjunctions of dependent outcomes as conditionals, estimates of $P(O_1 \& O_2)$ were more similar to estimates of $P(O_2/O_1)$ for dependent-outcome problems than for independent-outcome problems, and the two probabilities were identical for a higher proportion of subjects who judged the former problems.

Is it possible that the pattern of conjunction fallacies occurred, not because of the manipulation of conditional dependencies, but because of differences in the representativeness of dependent- and independent-outcome problems? For example, it could be argued that the win-lose conjunction was less representative for dependent than for independent outcomes (because it is less representative for the Oilers to lose the entire series than to lose one game), but that the lose-win conjunction was more representative for dependent outcomes (because it is more representative for the Oilers to win the entire series than to win one game). However, it is difficult to account for the closer correspondence of conjunctive and conditional probabilities for dependent-outcome than independent-outcome problems based on differences in representativeness. Furthermore, if differences in representativeness accounted for the pattern of the fallacy, then lose-lose conjunctions should have yielded a lower incidence of the fallacy for dependent than for independent outcomes (since losing both the first game and the series would be most unrepresentative of all) and this was not the case.

The finding that positive dependencies increase the incidence of the conjunction fallacy replicates the findings of previous research (e.g., Tversky and Kahneman, 1983). It has generally been assumed, however, that negative dependencies would reduce the incidence of the fallacy. For example, Tversky and Kahneman (1983) say "the conjunction fallacy was common in the Linda problem despite the fact that the stereotypes of bank teller and feminist are mildly incompatible" (p. 305). Yet the results of this study show that negative as well as positive dependencies can lead to conjunction fallacies as a result of conditional interpretations. Thus, these results suggest that the high incidence of the fallacy in the Bjorn Borg problem and similar problems used in other research (e.g., Yates & Carlson, 1986) is probably attributable to conditional interpretations of conjunctions rather than to reliance on representativeness.

The conjunctions in this study entailed an explicit temporal order of the component outcomes. It is an open question to what extent people treat conjunctions of dependent events as conditionals when the events do not have an explicit temporal order. This question is of importance because manipulating the representativeness relations between a target description and the component events of a conjunction commonly affects the conditional dependencies between the events. Most previous research has used conjunctions of a representative (R) and an unrepresentative (U) event, which tend to be negatively dependent (as shown in Experiment 1). Conditional interpretations of such conjunctions could either increase or decrease the incidence of the conjunction fallacy depending on (a) whether subjects assess P(R/U) or P(U/R), and (b) the strength of the dependency between R and U. The fallacy is most likely if subjects assess P(R/U), because this will often fall between P(R) and P(U). If, however, the negative dependency is very strong, then P(R/U) will fall below P(U) and the fallacy will not occur.

The incidence of the conjunction fallacy was also strongly influenced by the combination of component-outcome probabilities. Conjunctions of a high- and a low-probability outcome yielded a much higher incidence and magnitude of the fallacy than did conjunctions of high-probability outcomes or of low-probability outcomes. The low incidence of the fallacy for conjunctions of high-probability (or representative) outcomes does not correspond with what one would expect if subjects were judging conjunctions according to representativeness (since conjunctions of representative outcomes should be more representative than either of the outcomes). However, this finding is difficult to interpret because there was a pronounced ceiling effect for the subjective probabilities of the high-probability outcomes. In fact, 15% of subjects judged the probabilities of both component outcomes as 100%. Thus, the low incidence and magnitude of the fallacy for these conjunctions may have been partly an artifact of this ceiling effect.

Experiment 3

Incorrect Combination Rules and the Conjunction Fallacy

I have suggested that people may judge the probability of a conjunction indirectly, by combining the probabilities of its component events, even when they directly judge the component-event probabilities according to representativeness. This experiment examined to what extent the conjunction fallacy results from representativeness-mediated probability judgments or simply from the use of incorrect operations to combine component-event probabilities.

According to the incorrect-combination-rule interpretation, representativeness is related to the conjunction fallacy only because it influences subjects' estimates of the component-event probabilities. If so, subjects who make representativeness-mediated judgments and subjects who cannot rely on representativeness but are simply given similar combinations of component-event probabilities and asked to judge the conjunctive probability should commit the conjunction fallacy equally often. If, on the other hand, the conjunction fallacy stems from reliance on representativeness, then the latter group of subjects should be much less likely to commit the conjunction fallacy.

In this study, representativeness subjects were given problems (such as the Linda problem) for which they could base their probability judgments on representativeness. The conjunctions were composed of two representative (high-probability) events, two unrepresentative (low-probability) events, or one representative and one unrepresentative (one high- and one low-probability) event. Probability-combination subjects were given problems for which they could not make representativeness-mediated judgments, but could judge the conjunctive probability only by combining the component-event probabilities (which were simply provided in the problem). The conjunctions were composed of two high-probability events, two low-probability events, or one high- and one low-probability event. Thus, both groups worked with similar combinations of probabilities, but only the representativeness subjects could judge the conjunction directly according to representativeness.

If the conjunction fallacy is due to reliance on representativeness, probability-combination subjects should be much less likely than representativeness subjects to commit the conjunction fallacy. If the fallacy is entirely due to inappropriate combination rules, then both groups should commit the fallacy equally often and with the same combinations of component-event probabilities. In short, probability-combination subjects will show how much of the conjunction fallacy is due to incorrect combination-rules, whereas any increase in the fallacy for representativeness subjects will show how much it is due to other factors, such as reliance on representativeness.

Method

Subjects and Design

The sample was composed of 212 male and female introductory psychology students, who participated in partial fulfillment of a course requirement. The design was a 2 X 3 (Representativeness vs. Probability-Combination X Component-Event Probabilities) between-subjects factorial design. From 33 to 37 subjects were randomly assigned to each condition.

Stimuli

Representativeness problems. Two conjunction problems were adapted from the Linda and Bill problems used by Tversky and Kahneman (1983). These were selected because they have been used in several studies of the conjunction fallacy and are commonly presented as exemplars of problems evoking conjunction fallacies from representativeness-mediated probability judgments. The representativeness relations between the target descriptions and the component events were manipulated so that the conjunctions were composed of two high-probability events (HH), two low-probability events (LL), or one high- and one low-probability event (HL). These problems are presented below. The high-probability version of each event is listed first and the low-probability version is listed second. The HL version of the Linda problem contained the low-probability event A and the high-probability event B. The HL version of the Bill problem contained the high-probability event A and the low-probability event B.

Linda Problem:

> Linda is in her early 30's, 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 anti-nuclear demonstrations. What is the probability of each of the following?
>
> A. Linda is an avid reader [a bank teller].
>
> B. Linda is active in the feminist movement [extremely fashion-conscious].
>
> C. Linda is both A and B.

Bill problem:

> Bill is in his early thirties. He is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities. What is the probability of each of the following?
>
> A. Bill is an accountant [an excellent athlete].
>
> B. Bill is politically conservative [plays jazz for a hobby].
>
> C. Bill both A and B.

Probability-combination problem. This problem identified an imaginary planet and the creatures populating the planet. The problem provided the probability that a creature would have feature A and the probability that a creature would have feature B. Subjects were required to judge the probability that a creature would have both A and B. Three versions of the problem contained two high-probability features (80% and 75%), one high- and one low-probability feature (80% and 20%), or two low-probability features (25% and 20%). The problem read as follows:

> Imagine that you are exploring the Planet Kropiton. Kropiton is inhabited by Gronks. The probability that a Gronk will have blue hair is __%. The probability that a Gronk will have three eyes is __%. What is the probability that the first Gronk you meet will have both blue hair and three eyes?

Strategy questionnaires. These were designed to determine (a) whether representativeness subjects judged conjunctive probabilities directly or by combining the component-event probabilities, and (b) what combination-rules subjects used to combine component-event probabilities. Appendix D contains a sample strategy questionnaire used for the HH version of the Bill problem. Subjects given representativeness problems first indicated whether they had judged the conjunctive probability directly without using the probabilities of its component events, or based on the component-event probabilities. Subjects who reported using the latter strategy then indicated whether or not they had used each of several strategies to combine P(A) and P(B). These strategies included the following: (a) averaging, (b) averaging and adjusting up or down, (c) adding, (d) selecting the lower probability, (e) selecting the higher probability, (f) multiplying, (g) multiplying and adjusting up or down, (h) selecting P(A/B), (i) selecting P(B/A), (j) combining P(A) and P(B/A) by averaging or multiplication and (k) combining P(B) and P(A/B) by averaging or multiplication. These strategies were selected based partly on examination of verbal protocols collected from pilot subjects, and partly on the findings of previous research (Abelson et al., 1987; Goldsmith, 1978; Wyer, 1976; Yates & Carlson, 1986).

The strategy questionnaire for probability-combination subjects was identical except that they were not asked if they had judged the conjunction directly or based on the component-event probabilities (because they could not use the former strategy), and the questionnaire did not include strategies that entailed assessing the conditional dependencies of events A and B (strategies h through k).

Procedure

Introductory instructions were as in previous experiments. Subjects received a booklet containing either the two representativeness problems or the probability-combination problem with corresponding strategy questionnaires, and an answer booklet. Half the representativeness subjects judged the Bill problem first and half judged the Linda problem first. On the page following each problem, subjects filled out a strategy questionnaire. Following the probability judgment task, subjects completed a background questionnaire. This questionnaire was identical to that used in Experiment 1, except subjects were not given the option of indicating that they did not know the correct answer for any questions. They were asked to make their best guess even if they were uncertain of the correct answer. This change was made to determine what combination rules subjects thought were most reasonable, even if they did not think they knew the correct rule.

Results

Table 6 displays the mean probabilities of conjunctions and component events for representativeness and probability-combination subjects. The component-event probabilities for representativeness subjects are subjective probabilities; the corresponding probabilities for probability-combination subjects were simply provided in the problems. As can be seen in this table, representativeness and probability-combination subjects worked with similar combinations of component event probabilities. This correspondence was very close for LL and HL conjunctions, and moderately close for HH conjunctions.

The magnitude and incidence of the conjunction fallacy. The magnitude of the fallacy was obtained as in Experiment 1. For representativeness subjects, the incidence of the fallacy was coded by assigning subjects a 0, .5, or 1, according to whether they committed the fallacy for neither, one, or both of the problems (note that the mean of these incidence scores is equal to the proportion of problems yielding the fallacy). Because probability-combination subjects each judged one problem, they were assigned either a 0 or a 1. The incidence measure therefore had a greater variance for probability-combination subjects, but the heterogeneity of variance was slight (s^2, = 1.46s^2).

Table 7 displays the mean magnitude and incidence of the conjunction fallacy for representativeness and probability-combination subjects. These measures were analysed in 2 X 3 (Representativeness

Table 6

Mean Probabilities of the Conjunction (A&B) and its Component Events (A, B) for Experiment 3

Component Probabilities	Representativeness									Probability-Combination		
	Bill			Linda			Combined					
	P(A)	P(B)	P(A&B)	P(A)	P(B)	P(A&B)	P(A)	P(B)	P(A&B)	P(A)	P(B)	P(A&B)
High-High	70.7	50.8	53.1	74.4	68.9	71.6	72.5	59.8	62.3	80.0	75.0	71.2
High-Low	78.8	19.3	22.0	74.4	28.5	36.8	76.6	23.9	29.4	80.0	20.0	29.8
Low-Low	21.2	27.1	9.0	29.6	26.4	16.0	25.4	26.8	12.5	25.0	20.0	12.5

Note. Representativeness subjects estimated P(A), P(B), and P(A&B), whereas Probability-combination subjects were given P(A) and P(B) and estimated P(A&B).

Table 7

Mean Magnitude and Incidence of the Conjunction Fallacy for Experiment 3

	Problem Type			
	Representativeness			Probability-Combination
Component Probabilities	Bill	Linda	Combined	
High-High	3.9 (.45)	8.1 (.58)	6.0 (.52)	-3.7 (.49)
High-Low	2.8 (.38)	9.7 (.56)	6.3 (.47)	9.6 (.46)
Low-Low	-5.7 (.12)	-2.2 (.35)	-3.9 (.24)	-7.5 (.22)

vs. Probability-combination X Component-event Probabilities: HH, HL, or LL) analyses of variance. Overall, representativeness subjects were no more likely to commit the conjunction fallacy than probability-combination subjects, (.41 vs. .39) $F(1, 206) < 1$, nor did the representativeness vs. probability-combination manipulation interact with the manipulation of component-event probabilities $F(2, 206) < 1$. Individual contrasts showed that the incidence of the fallacy for representativeness and probability-combination subjects did not significantly differ for any of the component-event probabilities, $ts(206) < 1$. The mean magnitude of the fallacy was somewhat greater for representativeness subjects (mean = 2.74) than for probability-combination subjects (mean = -.53), although this difference was only marginally significant, $F(1, 206) = 3.52$, $p = .06$. There was, however, a significant interaction between the representativeness vs. probability-combination manipulation and the manipulation of component-event probabilities on the magnitude of the fallacy, $F(2, 206) = 4.70$, $p < .01$. As can be seen in Table 7, HL and LL combinations resulted in similar magnitudes of the fallacy for representativeness and probability-combination subjects, but HH combinations led to a higher magnitude of the fallacy for representativeness subjects. Individual comparisons showed that the magnitude of the fallacy for the two groups significantly differed only for HH conjunctions, $t(206) = 3.20$, $p < .01$; HL and LL $ts(206) < 1.20$, $ps > .20$.

The incidence of the fallacy was significantly affected by the manipulation of component-event probabilities, $F(2, 206) = 8.24$, $p < .001$. As can be seen in Table 7, both representativeness and probability-combination subjects were least likely to commit the fallacy with LL conjunctions, and about equally likely to commit the fallacy with HL and HH conjunctions. Individual contrasts showed that LL conjunctions resulted in a lower incidence of fallacies than either HH or HL conjunctions, $ts(206) > 3.25$, $ps < .001$.

As in Experiment 1, there were no significant problem-order effects for representativeness subjects on either the magnitude or incidence of the fallacy, $ps > .30$.

Strategy reports. Most representativeness subjects reported that they did not judge the conjunctive probability directly, but instead based it on the probabilities of the component events. This was reported by 70% of subjects for the Linda problem and 81% for the Bill problem. These percentages were not significantly affected by the manipulation of component-event probabilities for either problem, ^{2}s (2) < 2.40, $ps > .30$. Interestingly, subjects who reported judging the conjunction directly were less likely to commit the conjunction fallacy than those who reported basing their judgments on the component-event probabilities, $r(99) = -.23$, $p < .05$. However, this relation held for HH conjunctions ($r = -.41$, $p < .05$) and LL conjunctions ($r = -.31$, $p < .10$), but not for HL conjunctions ($r = -.04$).

Representativeness subjects were assigned a 0, .5, or 1.0 for each strategy, according to whether they had used it for neither, one, or both of the problems. Probability-combination subjects were assigned either a 0 or a 1.0 for each strategy. Data from 14 representativeness subjects who reported judging the conjunction directly for both problems were not included in the analyses of strategy reports.

Table 8 shows the mean strategy reports of representativeness and probability-combination subjects. As can be seen in this table, representativeness and probability-combination subjects' strategy reports were very similar. The two groups significantly differed only in their use of the addition strategy, $F(1, 194) = 8.70$, $p < .01$. Otherwise, there were no significant differences between strategy reports for the two groups, $Fs(1,194) < 1.50$, $ps > .20$.

By far the most common strategies reported by subjects in all conditions were those involving averaging the probabilities of the component events. Strategies involving multiplication, the normatively appropriate combination rule, were used for only a minority of problems. Taking the lower component probability as the estimate of the conjunctive probability was reported for 17% of problems, but using the higher probability was reported for only 4% of problems. Interestingly, representativeness subjects reported using the conditional probability of one event given the other [i.e., P(A/B)] as the estimate of the conjunctive probability for 19% of problems.

vs. Probability-combination X Component-event Probabilities: HH, HL, or LL) analyses of variance. Overall, representativeness subjects were no more likely to commit the conjunction fallacy than probability-combination subjects, (.41 vs. .39) $F(1, 206) < 1$, nor did the representativeness vs. probability-combination manipulation interact with the manipulation of component-event probabilities $F(2, 206) < 1$. Individual contrasts showed that the incidence of the fallacy for representativeness and probability-combination subjects did not significantly differ for any of the component-event probabilities, $ts(206) < 1$. The mean magnitude of the fallacy was somewhat greater for representativeness subjects (mean = 2.74) than for probability-combination subjects (mean = -.53), although this difference was only marginally significant, $F(1, 206) = 3.52$, $p = .06$. There was, however, a significant interaction between the representativeness vs. probability-combination manipulation and the manipulation of component-event probabilities on the magnitude of the fallacy, $F(2, 206) = 4.70$, $p < .01$. As can be seen in Table 7, HL and LL combinations resulted in similar magnitudes of the fallacy for representativeness and probability-combination subjects, but HH combinations led to a higher magnitude of the fallacy for representativeness subjects. Individual comparisons showed that the magnitude of the fallacy for the two groups significantly differed only for HH conjunctions, $t(206) = 3.20$, $p < .01$; HL and LL $ts(206) < 1.20$, $ps > .20$.

The incidence of the fallacy was significantly affected by the manipulation of component-event probabilities, $F(2, 206) = 8.24$, $p < .001$. As can be seen in Table 7, both representativeness and probability-combination subjects were least likely to commit the fallacy with LL conjunctions, and about equally likely to commit the fallacy with HL and HH conjunctions. Individual contrasts showed that LL conjunctions resulted in a lower incidence of fallacies than either HH or HL conjunctions, $ts(206) > 3.25$, $ps < .001$.

As in Experiment 1, there were no significant problem-order effects for representativeness subjects on either the magnitude or incidence of the fallacy, $ts > .30$.

Strategy reports. Most representativeness subjects reported that they did not judge the conjunctive probability directly, but instead based it on the probabilities of the component events. This was reported by 70% of subjects for the Linda problem and 81% for the Bill problem. These percentages were not significantly affected by the manipulation of component-event probabilities for either problem, ^{2}s (2) $<$ 2.40, $ps > .30$. Interestingly, subjects who reported judging the conjunction directly were less likely to commit the conjunction fallacy than those who reported basing their judgments on the component-event probabilities, $r(99) = -.23$, $p < .05$. However, this relation held for HH conjunctions ($r = -.41$, $p < .05$) and LL conjunctions ($r = -.31$, $p < .10$), but not for HL conjunctions ($r = -.04$).

Representativeness subjects were assigned a 0, .5, or 1.0 for each strategy, according to whether they had used it for neither, one, or both of the problems. Probability-combination subjects were assigned either a 0 or a 1.0 for each strategy. Data from 14 representativeness subjects who reported judging the conjunction directly for both problems were not included in the analyses of strategy reports.

Table 8 shows the mean strategy reports of representativeness and probability-combination subjects. As can be seen in this table, representativeness and probability-combination subjects' strategy reports were very similar. The two groups significantly differed only in their use of the addition strategy, $F(1, 194) = 8.70$, $p < .01$. Otherwise, there were no significant differences between strategy reports for the two groups, $Fs(1,194) < 1.50$, $ps > .20$.

By far the most common strategies reported by subjects in all conditions were those involving averaging the probabilities of the component events. Strategies involving multiplication, the normatively appropriate combination rule, were used for only a minority of problems. Taking the lower component probability as the estimate of the conjunctive probability was reported for 17% of problems, but using the higher probability was reported for only 4% of problems. Interestingly, representativeness subjects reported using the conditional probability of one event given the other [i.e., P(A/B)] as the estimate of the conjunctive probability for 19% of problems.

Table 8

Mean Strategy Reports of Representativeness and Probability-combination Subjects for Experiment 3

	Representativeness				Probability-Combination			
	High-High	High-Low	Low-Low	Overall	High-High	High-Low	Low-Low	Overall
Averaging	.77	.46	.57	.61	.73	.56	.73	.68
and adjustment down	.25	.19	.38	.27	.27	.27	.41	.32
and adjustment up	.12	.06	.05	.08	.11	.00	.03	.05
Addition	.08	.06	.02	.05	.00	.00	.00	.00
Multiplication	.08	.13	.21	.14	.22	.19	.22	.21
and adjustment down	.05	.02	.00	.02	.08	.08	.03	.06
and adjustment up	.00	.00	.00	.00	.00	.00	.00	.00
Lower probability	.12	.19	.21	.17	.08	.22	.06	.12
Higher probability	.02	.02	.02	.02	.03	.05	.03	.04
P(B/A)	.30	.13	.13	.19				
Averaging P(A) and P(B/A)	.40	.20	.21	.26				
Multiplying P(A) and P(B/A)	.05	.04	.04	.04				

Note. Figures for averaging strategy include averaging-and-adjustment strategies. Figures for multiplication strategy include multiplication-and-adjustment strategies.

For the most part, subjects' strategy reports were similar for HH, HL and LL conjunctions. However, the manipulation of component-event probabilities had significant effects on subjects' use of the averaging strategy, $F(2, 193) = 3.94$, $p < .05$, and their use of the averaging-and-adjustment-down strategy, $F(2, 193) = 3.12$, $p < .05$. Individual comparisons showed that subjects were less likely to use an averaging strategy for HL conjunctions than for HH conjunctions, $p < .01$, or for LL conjunctions, $p < .10$, and more likely to use an averaging-and-adjustment-down strategy for LL than for HH conjunctions, $p < .10$, or for HL conjunctions, $p < .05$.

Were subjects' strategy reports related to the incidence of the conjunction fallacy? This was indeed the case. Table 9 displays correlations between strategy reports and the incidence of the conjunction fallacy. Correlations were calculated only for strategies that had been used for at least 5% of problems. As would be expected, subjects who reported averaging or adding the component probabilities were more likely to commit the conjunction fallacy than those who did not. This was not the case, however, for subjects who reported averaging and then adjusting their estimate downward. As would also be expected, subjects who reported multiplying the component probabilities or taking the lower probability as their estimate were less likely to commit the conjunction fallacy than others. Representativeness subjects who reported using as their estimate the conditional probability of one event given the other were more likely to commit the conjunction fallacy than subjects who did not use this strategy. Representativeness subjects who reported averaging the probability of one component event with the conditional probability of the other [i.e., averaging P(A) and P(B/A)] were also more likely to commit the fallacy than others.

Responses to the background questionnaire showed that only 31.9% of subjects knew (or guessed) the correct combination-rule for judging conjunctive probabilities; 40.5% thought that averaging was the correct rule, 21.0% thought that selecting the lower component probability was appropriate, and 6.2% thought that adding was the correct rule. Knowledge of the correct combination-rule was weakly related to the

Table 9

Correlations Between Strategy Use and the Conjunction Fallacy for Experiment 3

	Representativeness		Probability-Combination	
	Magnitude	Proportion	Magnitude	Proportion
Averaging	.20	.51 ***	.37 ***	.52 ***
Averaging and adjustment down	-.14	-.05	-.07	-.18
Averaging and adjustment up	.19	.24 *	-.08	.09
Addition	.51 ***	.31 **	--	--
Multiplication	-.41 ***	-.40 ***	-.35 ***	-.35 ***
Lower probability	-.16	-.29 **	-.02	-.17
Higher probability	.05	.11	.22 *	.14
P(B/A)	.35 **	.29 **		
Averaging P(B) and P(A/B)	.33 **	.33 **		

Note. Correlations were calculated only for strategies that were used for at least 5% of problems.

* $p < .05$
** $p < .01$
*** $p < .005$

incidence of the fallacy for both representativeness subjects, $r(99) = -.22$, $p < .05$, and probability-combination subjects, $r(109) = -.30$, $p < .01$. However, neither subjects' level of training in probability theory and statistics nor their self-reported knowledge of probability theory was significantly related to the incidence of the fallacy, $ps > .10$.

Discussion

The results of this experiment provide good evidence that the conjunction fallacy in representativeness problems stems from incorrect combination rules, and that it is related to representativeness only insofar as representativeness affects the subjective probabilities of a conjunction's component events. Subjects given representativeness problems and those given probability-combination problems (a) committed the conjunction fallacy equally frequently, (b) committed the fallacy when working with similar combinations of component-event probabilities, and (c) reported using similar strategies in judging the conjunctive probability.

A minority of representativeness subjects reported that they had judged the conjunctive probabilities directly, rather than based on the probabilities of the component events. Importantly, these subjects were less likely to commit the fallacy for HH and HL conjunctions than subjects who reported basing their judgments on the component-event probabilities, and equally likely to commit the fallacy for HL conjunctions. This finding further supports the contention that the fallacy stems from incorrect combination rules, rather than from direct judgments of conjunctions according to representativeness.

The only substantive difference between representativeness and probability-combination subjects occurred for conjunctions of high-probability events (HH). The <u>incidence</u> of the fallacy for HH conjunctions did not significantly differ for the two groups, but the mean <u>magnitude</u> of the fallacy was significantly higher for representativeness subjects. This difference is probably attributable to the less-than-perfect matching of the component-event probabilities of HH conjunctions for the two groups. The subjective probabilities of the events for representativeness subjects were somewhat lower than the probabilities that probability-combination subjects were given, so that any ceiling effect would be more pronounced for the latter group. Furthermore, the difference between the two component probabilities was substantially greater for representativeness subjects (mean difference = 12.7) than for probability-combination subjects (difference = 5.0). Recall that the most common strategies subjects reported (and the strategies most strongly related to the fallacy) involved averaging the the component-event probabilities. An averaging strategy would result in an increasing magnitude of the fallacy as the difference between the component probabilities increases.

Differences in the component-event probabilities may also account for the disparity between the results of this experiment and those of Experiment 2 (as well as those of Wells, 1986) for HH conjunctions. Recall that in Experiment 2, HH conjunctions yielded a lower incidence of the fallacy than did HL conjunctions. In Experiment 2, the subjective probabilities of the two high-probability outcomes were identical for 67% of subjects. An averaging strategy would not result in conjunction fallacies for these subjects. It is noteworthy that in Wells' study, the mean subjective probabilities of the component events for HH conjunctions were very similar for both problems subjects were presented (.84 and .89; .74 and .81).

The fallacy occurred much less frequently for LL conjunctions than for either HH or HL conjunctions. In part, this result might be an artifact of differences between the component probabilities for the three kinds of conjunctions. For representativeness subjects, the mean difference between these probabilities was lower for LL conjunctions than for HH or HL conjunctions; for probability-combination subjects, the difference for LL conjunctions was smaller than for HL conjunctions and equal to HH conjunctions. An averaging-and-adjustment-down strategy (which many subjects in all conditions reported using) would be more likely to bring the conjunctive probability below the lower component probability when the component probabilities are very similar. Thus, even if subjects in all conditions used the same judgmental strategies, the incidence of the fallacy would be expected to be lowest for LL conjunctions. However, it was also the case that subjects were somewhat more likely to report an averaging-and-adjustment-down strategy for LL conjunctions than for HH or HL conjunctions. The greater use of this strategy for LL conjunctions may at least partly explain why they yielded such a low incidence of the fallacy.

Although the correspondence between the results for representativeness- and probability-combination-subjects was very close, the two kinds of problems differed in several ways: (a) representativeness subjects estimated the component-event probabilities themselves and probability-combination subjects were simply given these probabilities, (b) although the component-event probabilities that the two groups worked with were similar, they were not identical, and (c) there was substantial variability in the probabilities of each component event for representativeness subjects, but not for probability-combination subjects. Perhaps representativeness subjects would have been more likely than probability-combination subjects to commit the fallacy had the problems not differed in these ways (although it is somewhat awkward to account for the correspondence of the results by invoking differences between conditions). Experiment 4 was conducted to address this concern and to further explore the incorrect-combination-rule interpretation of the fallacy.

Experiment 4

Further Examination of Incorrect Combination Rules in Conjunction Problems

The pattern of results in Experiment 3 suggests that the individual components of the conjunction are probably judged according to representativeness (how "stereotypic the event is of the description), but that the conjunction itself is judged by combining the probabilities of its components. Thus, for example, people may use their stereotype of feminists to judge the probability that Linda is a feminist and their stereotype of bank tellers to judge the probability that she is a bank teller. But the representativeness interpretation of the conjunction fallacy implies that people also have a stereotype of the conjunction (i.e., a feminist bank teller), whereas my interpretation is that no such stereotype exists. Nevertheless, it might be that people could construct a stereotype of a feminist bank teller or an athletic accountant or a politically conservative jazz player or even a blue-haired, three-eyed Gronk. This interpretation is not unequivocally ruled out because previous research has conditioned both component events on the same target description.

Consider however a conjunction of two events that are conditioned on different target descriptions. For example, suppose subjects read a description of Linda and a description of Bill and then judge the probability that *Linda is a bank teller and Bill is an accountant.* In such a problem, subjects can directly judge the probabilities of the component events according to representativeness (i.e., by assessing Linda's similarity to bank tellers and Bill's similarity to accountants). However, it would make no sense for subjects to combine the two similarities to arrive at an overall similarity, or to construct a stereotype of the conjunction, because each component event is related to only one target description. For example, it makes no sense to talk of how similar *Linda and Bill* are to *bank tellers and accountants.* Rather, Linda's similarity to the stereotype of bank tellers is assessed, and Bill's similarity to the stereotype of accountants is assessed. Thus, although representativeness can be used to judge the component probabilities, it cannot be used to directly judge the conjunctive probability. Indeed, the only reasonable strategy for judging the conjunction is to combine the component probabilities.

If conjunction fallacies occur because people judge conjunctive probabilities based on assessments of overall similarity, then conjunctions such as the above, which are conditioned on different descriptions, should reduce the incidence of the conjunction fallacy. If, on the other hand, conjunction fallacies occur because people separately judge the probability of each component event and then use incorrect rules to combine these probabilities, these conjunctions should yield the same incidence of the fallacy as conjunctions conditioned on one target description.

Yates and Carlson (1986) report the results of a manipulation very similar to that just described. They found that conjunctions of events that were generated by unrelated processes (e.g., Congress will avoid raising taxes in this election year & Bo Derek will be nominated for an Academy Award for her performance in her latest film) were only marginally less likely to result in conjunction fallacies than conjunctions of events generated by related processes (e.g., President Reagan will provide for increased military spending & Reagan will provide support for unwed mothers), and that this difference did not attain significance. Because subjects could not judge conjunctions of events generated by different processes directly, but only based on the component-event probabilities, this finding supports the argument that the conjunction fallacy stems from incorrect combination-rules. However, Yates and Carlson used entirely different problems for the two kinds of conjunctions, which may have varied in ways other than whether the component events involved related or unrelated processes.

In this experiment, some subjects were given <u>standard representativeness</u> problems which required them to judge conjunctions of two events that were conditioned on the same target description. A second group of subjects were given <u>mixed representativeness</u> problems which required them to judge conjunctions of two events that were conditioned on different target descriptions. Each standard and mixed problem contained two target descriptions and required subjects to judge two conjunctions of events. The target descriptions and component events were identical for the two kinds of problems; the only difference was in which events were paired in the conjunctions. Other subjects were given probability-combination problems in which they could judge the probability of conjunctions based only on the probabilities of the component events. Each probability-combination subject was yoked to one standard- or mixed-representativeness subject so that the component-event probabilities he or she was given were obtained from the yoked partner's estimates of these probabilities in a representativeness problem. Thus, representativeness subjects and probability-combination subjects worked with exactly the same component-event probabilities. If the conjunction fallacy occurs because people directly judge conjunctions according to representativeness, then standard representativeness problems should yield a higher incidence of the fallacy than either mixed representativeness or probability-combination problems. If, on the other hand, representativeness is related to the fallacy only insofar as it affects the probabilities of component events, then all three kinds of problems should yield a similar incidence of the fallacy.

Method

Subjects and Design

The sample was composed of 80 male and female introductory psychology students (20 per condition) who participated in partial fulfillment of a course requirement. The design was a 2 X 2 (Standard- vs. Mixed-Conjunctions X Representativeness- vs. Probability-combination) between-subjects factorial design. Subjects were randomly assigned to conditions with the constraint that data from one representativeness subject had to be available for each probability-combination subject.

Stimuli

Representativeness problems. Two problems (each requiring subjects to judge two conjunctions) were adapted from those used in

Experiments 1 and 3. These problems are in Appendix E. Each problem contained descriptions of two target persons. Following these were four events (a representative and an unrepresentative event conditioned on each description) and two conjunctions of these events. For standard representativeness problems, each conjunction was composed of the two events conditioned on the same description (e.g., Linda is a bank teller and Linda is active in the feminist movement). Mixed representativeness problems contained the same target descriptions and component events as standard problems. However, each conjunction was composed of one event conditioned on the first description and one event conditioned on the second description (e.g., Linda is a bank teller & Bill is an accountant).

Probability-combination problems. Four of these problems were constructed (see Appendix F). As in Experiment 3, each problem identified an imaginary planet and the creatures populating the planet. The problem provided the probability that a creature would have feature A and that it would have feature B; subjects judged the probability that a creature would have both A and B. Each probability-combination subject judged all four problems; the component-event probabilities came from his or her yoked partner's estimates of these probabilities for the representativeness problems. Thus, yoked pairs worked with identical component-event probabilities.

Procedure

Subjects participated in small groups. Introductory instructions were as in previous experiments. Representativeness subjects received a booklet containing two problems (each requiring them to judge two conjunctions) with corresponding strategy questionnaires. Probability-combination subjects received a booklet containing four problems and strategy questionnaires. The strategy questionnaires were as in Experiment 3, except strategies that had been used by less than 5% of subjects were omitted. Once subjects judged all the problems, they completed a background questionnaire as in Experiment 3.

Results

Table 10 displays mean subjective probabilities of conjunctions and their component events for standard- and mixed-representativeness subjects and their yoked probability-combination partners. Table 11 displays the mean magnitude and incidence of the conjunction fallacy for each of these groups. Representativeness subjects given standard problems and those given mixed problems did not significantly differ in either the incidence or magnitude of the fallacy, incidence $F(1, 38) < 1$; magnitude $F(1, 38) = 2.49$, $p > .10$. Because the data of representativeness and probability-combination subjects are not independent (i.e., yoked pairs worked with the same component-event probabilities), differences between these groups were assessed using paired difference tests. Mean differences were calculated between yoked pairs in both the incidence and the magnitude of the fallacy, and the analyses were performed on these difference scores. There were no significant differences between representativeness subjects and their yoked probability-combination partners in either the incidence or magnitude of the fallacy, paired difference $ts(39) < 1$. The Standard vs. Mixed Problem X Representativeness vs. Probability-combination interaction was assessed by comparing the mean differences between yoked pairs for the two groups. These differences did not significantly vary for standard and mixed problems, incidence $F(1, 38) < 1$; magnitude $F(1, 38) = 1.38$, $p > .20$.

Strategy reports. For each problem, between 65% and 95% of representativeness subjects reported judging the conjunctive probability based on the component-event probabilities (overall mean = 82.5%). Only 2 subjects reported judging the conjunctive probability directly (without using the component probabilities) for more than half the problems, and only one reported this for all four problems. Representativeness subjects given standard problems judged the conjunction directly for a higher proportion of problems (mean = .28) than those given mixed problems (mean = .08), $t(38) = 2.60$, $p < .05$. As was the case for HL conjunctions in Experiment 3, there was no significant relation between whether or not subjects judged

Table 10

Mean Probabilities of the Conjunction (R&U) and its Representative (R) and Unrepresentative (U) Component Events for Experiment 4

	Problem 1						Problem 2						Combined		
	R_1	U_1	R_2	U_2	$(R\&U)_1$	$(R\&U)_2$	R_1	U_1	R_2	U_2	$(R\&U)_1$	$(R\&U)_2$	R	U	(R&U)
Standard															
Representativeness	74.4	26.5	66.3	19.0	32.1	23.6	80.8	19.1	73.7	30.4	23.4	28.4	73.0	23.7	26.9
Probability-Combination					30.1	25.2					37.1	34.7			31.8
Mixed															
Representativeness	80.5	30.4	44.8	32.3	43.8	31.5	79.5	10.8	76.2	22.6	27.6	35.2	70.3	24.0	34.3
Probability-Combination					38.8	30.6					25.8	27.5			30.7

Note. The subscripts 1 and 2 indicate whether the event was conditioned on the first or second target description in a pair. For standard-representativeness subjects, $(R\&U)_1$ was composed of R_1 and U_1; $(R\&U)_2$ was composed of R_2 and U_2. For mixed-representativeness subjects, $(R\&U)_1$ was composed of R_1 and U_2; $(R\&U)_2$ was composed of R_2 and U_1. The component-event probabilities that probability-combination subjects were given were identical to those estimated by the corresponding group of representativeness subjects.

Table 11

Mean Magnitude and Incidence of the Conjunction Fallacy for Experiment 4

	Problem 1		Problem 2		Combined
	$(R\&U)_A$	$(R\&U)_B$	$(R\&U)_A$	$(R\&U)_B$	$(R\&U)$
Standard					
Representativeness	7.3 (.55)	4.6 (.35)	4.4 (.50)	-.1 (.35)	4.0 (.44)
Probability-Combination	5.3 (.35)	6.2 (.50)	18.1 (.60)	6.2 (.45)	8.9 (.48)
Mixed					
Representativeness	13.6 (.50)	6.4 (.35)	16.8 (.60)	12.6 (.55)	12.4 (.50)
Probability-Combination	8.7 (.35)	5.5 (.45)	15.0 (.45)	4.9 (.40)	8.5 (.41)

Note. For standard-representativeness subjects, $(R\&U)_A$ was composed of R_1 and U_1; $(R\&U)_B$ was composed of R_2 and U_2. For mixed-representativeness subjects, $(R\&U)_A$ was composed of R_1 and U_2; $(R\&U)_B$ was composed of R_2 and U_1.

conjunctions directly and the incidence of the conjunction fallacy, $r(38) = .11$, $p > .40$.

Strategy reports were coded as in Experiment 3, by assigning each subject a number from .00 to 1.00 for each strategy, depending on the proportion of problems for which the strategy was used. Problems for which subjects reported judging the conjunction directly were not included in the calculation. Data from one subject who reported judging the conjunction directly for all four problems were not included in the analyses of strategy reports.

Table 12 displays mean strategy reports for representativeness subjects and their yoked probability-combination partners. As can be seen in this table, subjects in all conditions reported using similar strategies. Representativeness subjects given standard problems and those given mixed problems did not significantly differ in their use of any strategy, $ts(37) < 2.02$, $ps > .05$. Comparisons of representativeness and probability-combination subjects' strategy reports were carried out using paired difference tests. There were no significant differences between the strategy reports of the two groups, $ts(38) < 1.60$, $ps > .10$.

Table 13 displays the correlations between strategy reports and the incidence of the conjunction fallacy for representativeness subjects and probability-combination subjects. As in Experiment 3, the averaging and the averaging-and-adjustment-up strategies were positively related to the incidence of the fallacy. The higher-probability strategy was positively related to the incidence of the fallacy for probability-combination subjects, but not for representativeness subjects. The multiplication strategy and the lower-probability strategies were negatively related to the incidence of the fallacy, although the latter relation was significant only for probability-combination subjects. Correlations involving other strategies did not attain significance.

Responses to the background questionnaire showed that a higher percentage of these subjects knew the correct combination-rule for judging conjunctive probabilities (50%) than in previous studies. As well, subjects had more training in probability theory and statistics

Table 12

Mean Strategy Reports of Representativeness and Probability-combination Subjects for Experiment 4

	Standard		Mixed	
	Repres	Prob-comb	Repres	Prob-comb
Averaging	.38	.48	.55	.50
and adjustment down	.18	.19	.27	.16
and adjustment up	.13	.06	.06	.10
Addition	.07	.04	.20	.08
Multiplication	.33	.30	.18	.09
Lower probability	.07	.19	.15	.29
P(B/A)	.17		.14	
Averaging P(A) and P(B/A)	.17		.19	
Multiplying P(A) and P(B/A)	.13		.13	

Note. Figures for averaging strategy include averaging-and-adjustment strategies.

Table 13

Correlations Between Strategy Reports and the Conjunction Fallacy for Experiment 4

	Representativeness	Probability-Combination
Averaging	.54 ***	.68 ***
Averaging and adjustment down	.09	.38 *
Averaging and adjustment up	.35 *	.40 *
Addition	.06	.30
Multiplication	-.52 ***	-.52 ***
Lower probability	-.20	-.38 *
Higher probability	.09	.38 *
P(B/A)	-.10	
Averaging P(B) and P(A/B)	.11	
Multiplying P(B) and P(A/B)	-.02	

Note. Correlations were calculated only for strategies that were used for at least 5% of problems.

* $p < .05$
** $p < .01$
*** $p < .005$

than in previous studies: 41.3% had taken at least one university-level course covering probability or statistics (compared to 29.3% in Experiment 1 and 21.4% in Experiment 3). Probably for this reason, background questionnaire responses were more strongly related to the incidence of the fallacy than in previous experiments. The number of correct answers to the probability questions was related to the fallacy for both representativeness subjects, $r(38) = -.58$, $p < .001$, and probability-combination subjects, $r(38) = -.36$, $p < .05$. The same was the case for subjects' self-ratings of knowledge of probability theory, representativeness $r = -.33$, $p < .05$; probability-combination $r = -.36$, $p < .05$.

Discussion

The results of this experiment provide compelling evidence that the conjunction fallacy occurs because people base conjunctive probabilities on incorrect combinations of the component probabilities. The correspondence of the results for standard- and mixed-representativeness subjects, as well as their yoked probability-combination partners, was striking.

Probability-combination subjects worked with exactly the same component-event probabilities as their representativeness partners. Despite the fact that the former subjects could not judge either the conjunctive probability or the component-event probabilities according to representativeness, the incidence of the fallacy for the two groups was almost identical (.47 vs. .44) and both groups overestimated the conjunctive probability to the same extent (8.20 vs. 8.73).

Most representativeness subjects reported that they had not judged conjunctions directly, but had based their judgments on the probabilities of the component events. Consistent with the argument set forth in the introduction, subjects were somewhat more likely to report judging the conjunction directly for standard problems than for mixed problems. Importantly, however, the incidence of the fallacy did not significantly differ for the two kinds of problems. If anything, the incidence and magnitude of the fallacy were slightly greater for mixed problems than for standard problems. Because subjects given mixed problems could directly assess the similarity between the target descriptions and the conjunction, this finding provides further evidence that representativeness does not contribute to the conjunction fallacy independently of its effect on the component probabilities.

The mixed conjunction problems have an additional feature of interest from the present standpoint: the component events are clearly independent. For example, the probability that Jane is a mother is obviously unaffected by whether or not Bill is an accountant. Because mixed and standard problems yielded a very similar incidence of the fallacy, the results suggest that consideration of conditional dependencies between component events had no consistent effect on probability judgments for standard problems.

As in Experiment 3, the most common combination strategies reported by both representativeness and probability-combination subjects involved averaging the component probabilities, and these strategies were strongly implicated in the conjunction fallacy for both groups. The multiplication strategy was used somewhat more commonly than in Experiment 3, probably because of the higher level of statistical sophistication of subjects in this study. As in Experiment 3, subjects who relied on this strategy were much less likely to commit the fallacy than those who did not.

General Discussion

This research examined two alternatives to the representativeness explanation of the conjunction fallacy: It was proposed that the fallacy occurs (a) because people's natural representations of information and habitual reasoning strategies lead them to form inappropriate sample spaces in judging conjunctions, and (b) because people judge conjunctions by combining the probabilities of their component events using incorrect rules.

Inappropriate sample spaces. Two hypotheses regarding the contribution of inappropriate sample spaces to the fallacy were evaluated: (1) People preferentially form sample spaces from categories, so that problems requiring judgments of the form P(Category A & Category B/Target Description) lead to the fallacy because people sample category members to determine what proportion fit the target description. (2) People tend to treat conjunctions of dependent events as conditionals, and such a tendency can yield the conjunction fallacy both for positive dependencies (i.e., when one event facilitates the other) and for negative dependencies (i.e., when one event inhibits the other).

Although Experiment 1 supported the hypothesis that people preferentially form sample spaces from categories, there was no evidence that such a sampling preference contributed to the conjunction fallacy. The results do indicate, however, that it may be worthwhile to explore the contribution of inappropriate sample spaces to other fallacies of probabilistic or logical reasoning. Of particular interest is the finding that subjects who relied most heavily on inappropriate sampling strategies were least likely to consider the prior probabilities of the target events in evaluating their posterior probabilities (given the target description). This finding suggests that inappropriate sample spaces may be involved in the base-rate fallacy, which is characterized by neglect of prior probabilities in making probabilistic predictions. Many base-rate problems used in previous research (e.g., Ginosar & Trope, 1980; Kahneman & Tversky, 1973) have an instance-category structure:

subjects read a description of a target instance and judge the probability that the instance belongs to various categories differing in their base-rate frequency. This problem structure was most likely to evoke inappropriate sampling strategies, and, accordingly, resulted in least attention to prior probabilities.

An interesting finding to emerge from Experiment 1 is that subjects relied much less on the two sampling strategies and on representativeness in judging the probability of conjunctions than in judging the probabilities of component events. Thus, subjects must have used some other strategy to judge conjunctions. This finding is consistent with the results of Experiments 3 and 4, which showed that people often judge conjunctions indirectly, by combining their component-event probabilities.

Experiment 2 supported the hypothesis that people treat conjunctions of events with an explicit temporal order as conditionals when the component events are dependent, but not when they are independent. As in previous research, positive dependencies were found to increase the incidence of the conjunction fallacy. Negative dependencies increased the incidence of the fallacy for conjunctions of a low probability event followed by a high probability event but decreased it for conjunctions of a high probability event followed by a low probability event. The findings of Experiment 2 suggest that some conjunction fallacies attributed to reliance on representativeness stem largely from conditional interpretations of conjunctions.

Incorrect combination-rules. Experiments 3 and 4 provided strong evidence that the fallacy occurs primarily because people judge conjunctions by combining their component-event probabilities using incorrect rules. Representativeness probably is involved in the process by which people estimate the component-event probabilities, but there was no evidence (either in these experiments or in Experiment 1) that direct judgments of conjunctions according to representativeness contribute to the conjunction fallacy. The results of these two experiments can be summarized as follows: (1) Subjects who judged conjunctions following assessments of representativeness and subjects who could judge conjunctions only by combining their component-event probabilities committed the fallacy equally. Manipulating the (subjective or objective) probabilities of component events had parallel effects on the incidence of the fallacy for the two groups. (2) The majority of subjects who judged problems allowing assessments of representativeness reported that they had not judged conjunctions directly, but by combining the probabilities of their component events. Furthermore, subjects who reported judging conjunctions directly were no more likely than those who did not to commit the fallacy for conjunctions of a representative (R) and an unrepresentative (U) event (Experiments 3 and 4) and were less likely to commit the fallacy for RR and UU conjunctions (Experiment 3). (3) Subjects in all conditions reported similar judgment strategies and their strategy reports were predictably related to the incidence of the fallacy.

Some caution must be used in interpreting the results of the strategy questionnaires. Although very few subjects indicated that they had used other strategies, it is possible than none of the strategies included on the questionnaires exactly captured those that subjects used. Nevertheless, subjects' reports suggest that an averaging or averaging-and adjustment rule comes closest to capturing many subjects' combination strategies, and that the conjunction fallacy most commonly results from such a strategy. Subjects did, however, report using a wide variety of strategies, some of which were likely to yield the fallacy (e.g., adding) and some of which were not (e.g., lower probability, multiplication).

Is it possible that subjects really did not know what strategies they used, but inferred them from the relations among their probability estimates (cf. Nisbett & Wilson, 1977; Wilson & Stone, 1985)? For example, subjects whose estimates of the conjunctive probability fell between the two component probabilities may have inferred that they used an averaging rule, and those whose estimates fell below both component probabilities may have inferred that they used a multiplication rule. If so, then representativeness subjects and probability-combination subjects may have reported using similar strategies simply because they made similar probability estimates, even if they used different strategies. Recent research, however, has demonstrated that people do have privileged insight into their judgment strategies, and that such insight is not simply due to observing the relations between their responses and stimulus conditions (e.g., Gavanski & Hoffman, 1987; Kraut & Lewis, 1982). Furthermore, the self-report tasks in Experiments 3 and 4 simply required subjects to indicate whether or not they had attended to particular kinds of information and how they had combined two numbers to arrive at a third, and it seems implausible that people lack self-knowledge in such domains (cf. Ericsson & Simon, 1980).

At the very least, subjects' strategy reports indicate what strategies they thought were appropriate. To the extent that evaluation of subjective probabilities is under volitional control (i.e., that people are capable of formulating and executing a judgmental algorithm) one might expect people to rely on strategies they believe are appropriate, regardless of whether they are directly aware of the execution of the algorithm.

Even though probability-combination and mixed-representativeness subjects could not judge conjunctions according to representativeness, it might be argued that the rules they used to combine probabilities were derived from the rules they habitually use in combining similarities. There are, however, several reasons to doubt that probability-combination rules are <u>derived</u> from similarity-combination rules (although people may sometimes use the same rules to combine probabilities and to combine similarities). First, consider that the most frequent rule subjects reported was averaging. Research from the information integration tradition (e.g., Anderson, 1965, 1974; Shanteau, 1975; Troutman & Shanteau, 1977) has shown averaging to be a ubiquitous strategy of information integration, which is not restricted to judgments involving assessments of representativeness.

Second, the effects of representativeness manipulations on the incidence of the fallacy do not correspond with what would be expected if rules for combining probabilities are derived from rules for combining similarity. If similarity is based on assessments of common and distinctive features, then conjunctions of representative events should be more representative than either component event. The subjective probability of a representive conjunction should, therefore, be higher than both of its component probabilities (or at the very least, equal to the higher component probability). However, in Experiment 3, this was the case for only 19% of representative conjunctions yielding the fallacy (17% for the corresponding group of probability-combination subjects). Furthermore, conjunctions of representative events resulted in either fewer fallacies than conjunctions of a representative and an unrepresentative event (Experiment 2) or an equal incidence of the fallacy (Experiment 3).

Although a substantial proportion of subjects in all experiments committed the conjunction fallacy, the incidence of the fallacy in this research was substantially lower than that found in some previous studies. This was the case both for problems adapted from previous research and newly constructed problems. Is it possible that some experimental artifact lowered the incidence of the fallacy in this research? For example, perhaps judging several problems or filling out strategy reports made subjects less likely to commit the fallacy. However, there were no order effects for either of the two experiments in which problem order was varied: problems presented first (before subjects had seen the other problems or strategy questionnaires) yielded almost exactly the same incidence of the fallacy as problems presented last (after subjects had judged other problems and made strategy reports). Thus, the low incidence of the fallacy is not attributable to either of these factors.

Two methodological differences between this and previous research may account for the difference in the incidence of the fallacy. First, in many previous studies, subjects rank-ordered the probabilities of the conjunction and the component events whereas in this research, subjects estimated these probabilities. Morier & Borgida (1984) have shown that estimating probabilities reduces the incidence of the fallacy relative to rank-ordering probabilities. In fact, a rank-ordering procedure artificially increases the incidence of the fallacy. With a probability estimation procedure, a high proportion of people judge the conjunctive probability as equal to the lower component probability. Forcing them to chose one as higher if they have no preference will result in the conjunction fallacy for about half these subjects. A second methodological difference had to do with the wording of the conjunctions. All of the conjunctions in this research were prefaced with "both" (as in "both A and B"). The purpose of this change was to prevent subjects from treating the conjunction as a disjunction by making it explicitly clear that both component events must be true for the conjunction to be true. Morier and Borgida (1984) report that the incidence of the fallacy for the Linda problem was substantially reduced when subjects judged a disjunction (i.e., "Linda is a bank teller or is active in the feminist movement") as well as a conjunction. Clearly, including both a disjunction and a conjunction is similar to modifying the conjunction to read "both A and B," in that both procedures reduce the likelihood that subjects will treat the conjunction as a disjunction. I have also found in previous research (although this was not manipulated experimentally) that the same problems tend to yield a lower incidence of the fallacy when the conjunctions read "both A and B" than when they simply read "A and B".

All the experiments reported here entailed what Tversky and Kahneman (1983) call direct tests of the conjunction fallacy, because subjects judged both the conjunction and its component events. As Tversky and Kahneman point out, direct tests allow subjects to compare the probabilities of the conjunction and its components, perhaps making the inclusion relation between them more evident. Although it is likely that using direct tests resulted in a lower incidence of the fallacy than would have been the case with indirect tests, many studies interpreted as supporting the representativeness explanation of the fallacy have also used direct tests. More importantly, the same factors that may have reduced the incidence of the fallacy for subjects who judged standard representativeness problems would have had comparable effects on those who judged mixed representativeness problems or probability-combination problems. For example, subjects in all three groups could compare their estimates of the conjunctive probability with the component-event probabilities. Thus, there is no reason to assume that using direct tests had differential effects for the three groups (of course, indirect tests are not possible for probability-combination problems).

Perhaps judging the component-event probabilities before the conjunctive probability made subjects less likely to judge the conjunction directly according to representativeness, and more likely to base their judgments on the component probabilities, than would otherwise be the case. Even if this were the case, it seems unlikely that this procedure reduced the incidence of the fallacy, since subjects who judged the conjunction directly were no more likely to commit the fallacy than those who based their judgments on the component probabilities (and, in fact, were less likely to commit the fallacy for HH and LL conjunctions). Nevertheless, I collected some supplementary data to determine if people judge conjunctions directly when they do not first judge the component-event probabilities. A small sample of subjects (n = 12) judged the standard representativeness problems used in Experiment 4; however, they first judged the probability of the conjunction and then, on the following page, the component-event probabilities. These subjects reported basing their judgments on the component probabilities for 70.8% of the problems, and only one subject reported directly judging all four conjunctions. Thus, people commonly judge conjunctions of a representative and an unrepresentative event based on the probabilities of their component events, even when they have not previously estimated these probabilities.

Clearly, using incorrect rules to combine probabilities will result in less than optimal judgments of compound events. How serious are the consequences of such departures from optimality for everyday judgments and decisions? Surprisingly, the consequences may be trivial in many judgment situations. Suppose, for example, you are interviewing applicants for a job requiring someone who both gets along well with others and is extremely methodical. You may decide whom to hire based on a rank-ordering of applicants according to your subjective probability that they will meet both these criteria. Assuming independence, a normatively appropriate strategy for combining P(gets along well) and P(extremely methodical) is to multiply the two probabilities. However, many incorrect rules will yield a rank-ordering of conjunctive probabilities highly similar to that obtained through a normatively appropriate rule. The outputs from an averaging (or addition) and a multiplication rule, for example, are commonly highly correlated (Rorer, 1971; Wyer, 1976). Indeed, in Experiment 3, the correlation between the mean and the product of P(A) and P(B) was .86 for Problem 1 and .92 for Problem 2. Thus, as long as incorrect combination-rules are applied consistently, they may have few detrimental effects for the great many decisions that are based on a rank-ordering of conjunctive probabilities.

If, however, a decision depends on the value of a conjunctive probability rather than on a rank ordering of conjunctive probabilities, the consequences of incorrect combination-rules may be serious. Suppose, for example, you are deciding whether or not to purchase Stock A. Stock A will pay a dividend only if both a low probabili[illegible] and a high probability event (H) occur. Using an incorrect [illegible] as averaging to combine P(L) and P(H) may lead to a decis[illegible]ase the stock, which could be a very expensive decision.

Incorrect combination rules may have the most serious consequences for decisions that require comparing the probability of a conjunction with the probabilities of simple events (e.g., the conjunction's components). However, it is difficult to imagine many ecologically meaningful judgment situations involving such a comparison (although it is no doubt possible to construct situations of this kind).

This research provides further evidence that the rules governing subjective judgments of probability deviate sharply from those governing mathematical probability theory: statistical intuitions commonly violate the Kolmogorov axioms. It seems implausible, then, that these intuitions would correspond to the more complex probability theorems derived from these axioms. For example, one should not expect statistically naive people to be capable of revising probabilistic judgments in accordance with the dictates of Bayes' theorem.

Does this mean that people's intuitions about uncertainty are fundamentally flawed? This inference is not warranted. In the first place, there is little evidence that people naturally quantify uncertainty as subjective probability, or that they typically base important judgments and decisions on subjective probabilities. For example, people might naturally quantify uncertainty as subjective relative frequency. Such judgments have been shown to obey different rules than subjective probabilities and to correspond more closely with the dictates of probability theory (Tversky & Kahneman, 1983). Furthermore, it is unclear that people even translate their representations of uncertainty into numerical form in everyday judgment situations. Fallacies of probabilistic inference may arise during the translation process, or when people attempt to integrate numerically represented information about uncertainty. People may generally make decisions and choices based on non-numerical representations that, unlike subjective probabilities, are in accordance with the laws of probability theory. Until we know more about how people naturally represent uncertainty we cannot specify how poor or good they are at quantifying and responding appropriately to uncertainty in their natural environment.

Footnotes

[1] All t tests were calculated using the mean square error and degrees of freedom for the entire 6-group design.

[2] All correlations between strategy ratings and subjective probabilities were calculated using only ratings that applied to the particular probability judgment in question, i.e., P(R), P(U), or P(R & U).

[3] Separate analyses were carried out for each problem to avoid confounding of within- and between-subjects error. There were also no significant order effects for Problems 2 or 3, each of which served in either the second or the third position.

[4] Describing these outcomes as representative or unrepresentative is consistent with Tversky and Kahneman's (1983) treatment of outcomes in the Bjorn Borg problem. It should, however, be pointed out that these outcomes may be representative or unrepresentative precisely by virtue of their probabilities. Thus, it is questionable if representativeness can serve as an explanatory construct in this context.

[5] The component events were joined with the connective "and" rather than "but," which Tversky and Kahneman used for the Bjorn Borg problem. Yates and Carlson (1986) found this change to have no effect on the incidence of the conjunction fallacy.

[6] All t tests were calculated using the mean square error and degrees of freedom for the entire 6-group design.

[7] Because a directional pattern of means was predicted for the dependent- vs. independent-outcome conditions, one-tailed significance tests were used for these contrasts.

References

Abelson, R. P., Leddo, J., & Gross, P. H. (1987). The strength of conjunctive explanations. Personality and Social Psychology Bulletin, 13, 141-155.

Anderson, N. H. (1965). Averaging versus adding as a stimulus-combination rule in impression formation. Journal of Experimental Psychology, 70, 394-400.

Anderson, N. H. (1974). Information integration theory: A brief survey. In D. H. Krantz, R. C. Atkinson, R. D. Luce, and P. Suppes (Eds.), Contemporary Developments in Mathematical Psychology (Vol. 2). San Francisco: Freeman.

Beyth-Marom (1981). The subjective probability of conjunctions. Decision Research Report 81-12.

Braine, M. D. S. (1979). On the relation between the natural logic of reasoning and standard logic. Psychological Review, 85, 1-21.

Crandall, C. S., & Greenfield, B. (1986). Understanding the conjunction fallacy. A conjunction of effects? Social Cognition, 4, 408-419.

Cohen, J., & Hansel, C. E. M. (1957). The nature of decisions in gambling. Equivalence of single and compound subjective probabilities. Acta Psychologica, 13, 357-370.

Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky, (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press.

Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. Psychological Bulletin, 99, 3-19.

Ericsson, K. A., & Simon, H. A. (1980). Verbal reports as data. Psychological Review, 87, 215-254.

Falk, R., & Bar-Hillel, M. (1983). Probabilistic dependence between events. Two-Year College Mathematics Journal, 14, 240-247.

Fischhoff, B. (1985). Judging unlikely conjunctions. Decision Research Report 85-2.

Gavanski, I., & Hoffman, C. (1987). Awareness of influences on one's own judgments: The roles of covariation detection and attention to the judgment process. Journal of Personality and Social Psychology, 52, 453-463.

Goldsmith, R. W. (1978). Assessing probabilities of compound events in a judicial context. Scandinavian Journal of Psychology, 19, 103-110.

Ginosar, Z., & Trope, Y. (1980). The effects of base rates and individuating information on judgments about another person. Journal of Experimental Social Psychology, 16, 228-242.

Hays, W. L. (1981). Statistics (3rd Ed.). New York: Holt, Rinehart and Winston.

Kahneman, D., Slovic, P., & Tversky, A. (1982). Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press.

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. Cognitive Psychology, 3, 430-454.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. Psychological Review, 80, 237-251.

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

Kraut, R. E., & Lewis, S. H. (1982). Person perception and self-awareness: Knowledge of influences on one's own judgments. Journal of Personality and Social Psychology, 42, 448-460.

Lichtenstein, Fischoff, & Philips, (1982). Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, & A. Tversky, (Eds.), Judgment under uncertainty: Heuristics and biases. New York: Cambridge University Press.

Markus, H., & Zajonc, R. B. (1985). The cognitive perspective in social psychology. In G. Lindzey & E. Aronson (Eds.), Handbook of social psychology, 3rd Edition. Reading, MA: Addison-Wesley.

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

Nisbett, R. E., & Ross, L. (1980). Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, N.J.: Prentice-Hall.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. Psychological Review, 84, 231-259.

Rorer, L. G. (1971). A circuitous route to bootstrapping selection procedures. In H. B. Haley, A. G. D'Costa, & A. M. Schafer (Eds.), Personality Measurement in Medical Education. Des Plaines, Ill.: Association of American Medical Colleges.

Shanteau, J. (1975). Averaging versus multiplying combination rules in inference judgment. Acta Psychologica, 39, 83-89.

Slovic, P. (1969a). Differential effects of real vs. hypothetical payoffs upon choices among gambles. Journal of Experimental Psychology, 80, 434-437.

Slovic, P. (1969b). Manipulating the attractiveness of a gamble without changing its expected value. Journal of Experimental Psychology, 79, 139-145.

Troutman, C. M., & Shanteau, J. (1977). Inferences based on nondiagnostic information. Organizational Behavior and Human Performance, 19, 43-55.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 5, 207-232.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. Science, 185, 1124-1131.

Tversky, & Kahneman, (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), Progress in Social Psychology. Hillsdale, N.J.: Erlbaum.

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

Appendix A

Conjunction Problems for Experiment 1

Instance-Category Problems

1. Linda is in her early 30's, 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 anti-nuclear demonstrations. What is the probability of each of the following?

A. Linda is a bank teller.
B. Linda is active in the feminist movement.
C. Linda is both a bank teller and active in the feminist movement.

2. Jason is in his late 20's, single and lives alone. He has travelled extensively and always enjoys seeing new things. He likes to go to clubs to hear new bands, he writes poetry and plays the guitar. What is the probability of each of the following?

A. Jason is an artist.
B. Jason is a computer programmer.
C. Jason is both an artist and a computer programmer.

3. Jane is in her late 20's and divorced. She is assertive and outspoken. She is university-educated, well-dressed and works for a large successful corporation. Her career requires that she travel frequently. What is the probability of each of the following?

A. Jane is an executive.
B. Jane is a mother.
C. Jane is both an executive and a mother.

4. Tracy is in her early twenties. She is not married and has a pre-school child. She works and she lives in a small apartment and spends most of her free time watching T.V. or playing with her son. What is the probability of each of the following?

A. Tracy is a waitress.
B. Tracy is a university graduate.
C. Tracy is both a waitress and a university graduate.

Instance-Instance Problems

1. Linda is in her early 30's, 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 anti-nuclear demonstrations. What is the probability of each of the following?

   A. Linda reads harlequin romances in her spare time.
   B. Linda has participated in a pro-abortion demonstration.
   C. Linda both reads harlequin romances in her spare time and has participated in a pro-abortion demonstration.

2. Jason is in his late 20's, single and lives alone. He has travelled extensively and always enjoys seeing new things. He likes to go to clubs to hear new bands, he writes poetry and plays the guitar. What is the probability of each of the following?

   A. Jason has unusual taste in clothing.
   B. Jason wants to marry a traditional woman.
   C. Jason both has unusual taste in clothing and wants to marry a traditional woman.

3. Jane is in her late 20's and divorced. She is assertive and outspoken. She is university-educated, well-dressed and works for a large successful corporation. Her career requires that she travel frequently. What is the probability of each of the following?

   A. Jane puts her career ahead of her personal life.
   B. Jane does needlepoint for a hobby.
   C. Jane both puts her career ahead of her personal life and does needlepoint for a hobby.

4. Tracy is in her early twenties. She is not married and has a pre-school child. She works and she lives in a small apartment and spends most of her free time watching T.V. or playing with her son. What is the probability of each of the following?

   A. Tracy would like to get married.
   B. Tracy enjoys going to the symphony.
   C. Tracy both would like to get married and enjoys going to the symphony.

Category-Instance Problems

1. Linda is a feminist. What is the probability of each of the following?

   A. Linda reads harlequin romances in her spare time.
   B. Linda has participated in a pro-abortion demonstration.
   C. Linda both reads harlequin romances in her spare time and has participated in a pro-abortion demonstration.

2. Marie is a Roman Catholic nun. What is the probability of each of the following?

   A. Marie enjoys spending time alone with her thoughts.
   B. Marie spends much of her free time watching T.V..
   C. Marie both enjoys spending time alone with her thoughts and spends much of her free time watching T.V..

3. Margaret is a middle-aged housewife. What is the probability of each of the following?

   A. Margaret regularly watches soap operas.
   B. Margaret likes rock music.
   C. Margaret both regularly watches soap operas and likes rock music.

4. Kirk is a middle-aged playboy. What is the probability of each of the following?

   A. Kirk is very concerned about becoming less attractive as he gets older.
   B. Kirk donates a lot of his money to charity.
   C. Kirk both is very concerned about becoming less attractive as he gets older and donates a lot of his money to charity.

Appendix B

Sample Strategy Questionnaire for Experiment 1

Following are several strategies you may have used in judging the probability that Linda is a bank teller. For each strategy, please use the rating scale on the answer sheet to indicate how much you relied on it. If you did not rely on the strategy at all, then circle the 0 on the rating scale. If you did rely on a particular strategy, then use the rating scale to indicate how heavily you relied on it (1 = used this strategy a little; 7 = used this strategy very heavily).

1.1 *I estimated what proportion of people who fit Linda's description would be bank tellers.*

1.2 *I estimated what proportion of bank tellers would fit Linda's description.*

1.3 *I estimated what proportion of females would be bank tellers.*

1.4 *I assessed the similarity between Linda's description and my stereotype of bank tellers.*

1.5 *If you did not rely on any of the above strategies, or also relied on some other strategy, then please indicate in the blank space on the answer sheet what other strategy or strategies you used.*

Please use the rating scales on the answer sheet to indicate how much you relied on each of the following strategies in judging the probability that Linda is active in the feminist movement.

1.6 *I estimated what proportion of people who fit Linda's description would be feminists.*

1.7 *I estimated what proportion of feminists would fit Linda's description*

1.8 *I estimated what proportion of females would be feminists.*

1.9 *I assessed the similarity between Linda's description and my stereotype of feminists.*

1.10 *If you did not rely on any of the above, or also relied on some other strategy, then please indicate in the blank space on the answer sheet what other strategy or strategies you used.*

Please use the rating scales on the answer sheet to indicate how much you relied on each of the following strategies in judging the probability that Linda is both a bank teller and active in the feminist movement.

1.11 *I estimated what proportion of people who fit Linda's description would be bank tellers and feminists.*

1.12 *I estimated what proportion of bank tellers and feminists would fit Linda's description.*

1.13 *I estimated what proportion of females would be bank tellers and feminists.*

1.14 *I assessed the similarity between Linda's description and my stereotypes of bank tellers and feminists.*

1.15 *If you did not rely on any of the above, or also relied on some other strategy, then please indicate in the blank space on the answer sheet what other strategy or strategies you used.*

Appendix C

Background Questionnaire for Experiment 1

INSTRUCTIONS: The following questions are to inform us how much training and background you have in probability theory or statistics. Please do not guess if you don't have a good idea what the correct answer is. Instead, select the last alternative, which indicates you don't know the answer.

1. Consider two events, A and B. The two events are independent (in other words, the probability of one occurring isn't affected by whether or not the other one has occurred). What is the appropriate rule to find the probability that both A and B will occur?

   a. take the lower of the two events' probabilities
   b. take the higher of the two events' probabilities
   c. multiply the probability of A by the probability of B
   d. average the probabilities of A and B
   e. add the probabilities of A and B
   f. there is insufficient information to determine the appropriate rule
   g. I don't know what rule should be used

2. If two events, A and B, are mutually exclusive, then

   a. they are independent
   b. if one has occurred, the other can't occur
   c. if one has occurred, the other must occur
   d. they must be equally probable
   e. I don't know

3. If two events, A and B, are mutually exclusive, what rule should be used to find the probability that either A or B will occur?

   a. take the lower of the two event's probabilities
   b. take the higher of the two events' probabilities
   c. multiply the probability of A by the probability of B
   d. average the probabilities of A and B
   e. add the probabilities of A and B
   f. there is insufficient information to determine the appropriate rule
   g. I don't know what rule should be used

4. How much academic background do you have in probability theory and statistics?

   a. none or extremely little
   b. I studied probability or statistics in high school
   c. I have had only one university or college course covering probability and/or statistics
   d. I have had more than one university or college course covering probability and/or statistics.

5. How would you evaluate your overall understanding of probability theory and the rules of probability? Please use the rating scale on the answer sheet.

Appendix D

Sample Strategy Questionnaire for Experiment 3

1.01 In judging the probability that Bill is both an accountant and is politically conservative, which of the following did you do?

A. I separately estimated the probability that Bill is an accountant and the probability that Bill is politically conservative and then based my estimate that he is both on these two probabiPities.

B. I directly estimated the probability that Bill is both an accountant and politically conservative without using the individual probabilities.

If you used A above, how did you use the two probabilities to arrive at your estimate? Please indicate on the answer sheet whether or not you used each of the following strategies. Please read over all the strategies before indicating which you used. Do not feel obliged to say that you used any of these strategies if you didn't. Many strategies you may have used are not listed here.

1.1 I roughly averaged the probability that Bill is an accountant and the probability that Bill is politically conservative to get my estimate.

1.2 I roughly averaged the two probabilities and then adjusted my estimate upward or downward (please indicate which on the answer sheet).

1.3 I roughly added the two probabilities to get my estimate.

1.4 I roughly multiplied the two probabilities to get my estimate.

1.5 I roughly multiplied the two probabilities and then adjusted my estimate upward or downward (please indicate which on the answer sheet).

1.6 I took the lower of the two probabilities as my estimate.

1.7 I took the higher of the two probabilities as my estimate.

1.8 I judged the probability that Bill will be an accountant if he is politically conservative and used that as my estimate.

1.9 I judged the probability that Bill will be politically conservative if he is an accountant and used that as my estimate.

1.10 I roughly averaged the probability that Bill is an accountant with the probability that Bill will be politically conservative if he is an accountant.

1.11 I roughly multiplied the probability that Bill is an accountant with the probability that Bill will be politically conservative if he is an accountant.

1.12 I roughly averaged the probability that Bill is politically conservative with the probability that Bill will be an accountant if he is politically conservative.

1.13 I roughly multiplied the probability that Bill is politically conservative with the probability that Bill will be an accountant if he is politically conservative.

1.14 If you did not rely on any of the above, or also relied on some other strategy (or strategies), then please indicate in the blank space on the answer sheet what other strategy you used.

Appendix E

Representativeness Problems for Experiment 4

Problem 1

Following are descriptions of two people, Jane and Bill.

Jane is in her late 20's and divorced. She is assertive and outspoken. She is university educated, well dressed and works for a large successful corporation. Her career requires that she travel frequently.

Bill is intelligent, but unimaginative, compulsive, and generally lifeless. In school, he was strong in mathematics but weak in social studies and humanities.

What is the probability of each of the following?

A. Jane is an executive.
B. Jane is a mother.
C. Bill is an accountant.
D. Bill plays jazz for a hobby.

Conjunctions for Standard Problem:

E. Jane is both an executive and a mother.
F. Bill both is an accountant and plays jazz for a hobby.

Conjunctions for Mixed Problem:

E. Both Jane is an executive and Bill plays jazz for a hobby.
F. Both Jane is a mother and Bill is an accountant.

Problem 2

Following are descriptions of two people, Linda and Jason.

Linda is in her early 20's, 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 anti-nuclear demonstrations.

Jason is in his late 20's, single and lives alone. He has travelled extensively and always enjoys seeing new things. He likes to go to clubs to hear new bands, he writes poetry and plays the guitar.

What is the probability of each of the following?

- *A. Linda is a bank teller.*
- *B. Linda is active in the feminist movement.*
- *C. Jason is a computer programmer.*
- *D. Jason is an artist.*

Conjunctions for Standard Problem:

- *E. Linda is both a bank teller and active in the feminist movement.*
- *F. Jason is both a computer programmer and an artist.*

Conjunctions for Mixed Problem:

- *E. Both Linda is a bank teller and Jason is an artist.*
- *F. Both Linda is active in the feminist movement and Jason is a computer programmer.*

Appendix F

Probability-combination Problems for Experiment 4

1. Imagine that you are exploring the planet Kropiton. Kropiton is inhabited by Gronks. The probability that a Gronk will have blue hair is __%. The probability that a Gronk will have three eyes is __%. What is the probability that the first Gronk you meet will have both blue hair and three eyes?

2. Imagine that you are exploring the planet Zilcon. Zilcon is inhabited by Skeexes. The probability that a Skeex will eat grass is __%. The probability that a Skeex will drool a lot is __%. What is the probability that the first Skeex you meet will both eat grass and drool a lot?

3. Imagine that you are exploring the planet Fronkus. Fronkus is inhabited by Yumps. The probability that a Yump will eat humans is __%. The probability that a Yump walks backward is __%. What is the probability that the first Yump you meet will both eat humans and walk backward?

4. Imagine that you are exploring the planet Meenus. Meenus is inhabited by Igys. The probability that an Igy will have fuzzy ears is __%. The probability that an Igy will be bad-tempered is __%. What is the probability that the first Igy you meet will both have fuzzy ears and be bad-tempered?