

42

How Alike Is It Versus How Likely Is It: A Disjunction Fallacy in Probability Judgments

Maya Bar-Hillel and Efrat Neter

One event cannot be more probable than another that includes it. Judging $P(A \& B)$ to be higher than $P(A)$ has been called the *conjunction fallacy*. This study examined a *disjunction fallacy*. Ss received brief case descriptions and ordered 7 categories according to 1 of 4 criteria: (a) probability of membership, (b) willingness to bet on membership, (c) inclination to predict membership, and (d) suitability for membership. The list included nested pairs of categories (e.g., Brazil-South America). Ranking a category more probable than its superordinate, or betting on it rather than its superordinate, is fallacious. Prediction, however, may be guided by maximizing informativeness, and suitability need conform to no formal rule. Hence, for these 2 criteria, such a ranking pattern is not fallacious. Yet ranking of categories higher than their superordinates was equally common on all 4 criteria. The results support representativeness against alternative interpretations.

The extension rule in probability theory states that if A is a subset of B, then the probability of A cannot exceed that of B. A special case of the extension rule is the conjunction rule, which states that the probability of A&B can exceed the probability of neither A nor B, since it is contained in both.

Tversky and Kahneman (1983) demonstrated that, under certain circumstances, people predictably and systematically violate the conjunction rule. In one study, they gave subjects the following description:

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. (p. 297)

This was followed by a list of eight possible outcomes, each describing possible activities of Linda at the present time (her job, her interests, or both). Subjects were asked to rank order the outcomes by the probability that they describe Linda's current activities. Of the eight, one was representative of Linda ("Linda is active in the feminist movement"), one was unrepresentative of Linda ("Linda is a bank teller"), and one was a conjunction of these two ("Linda is a bank teller and is active in the feminist movement"). A large majority of the subjects (85%) rated the conjunctive outcome, "Linda is a bank teller and is active in the feminist movement," more probable than "Linda is a bank teller."

This result was predicted from the representativeness hypothesis: "Representativeness is an assessment of the degree of correspondence between a sample and a population, an instance and a category, an act and an actor or, more generally, between

an outcome and a model" (Tversky & Kahneman, 1983, p. 295). Kahneman and Tversky (1972, 1973) provided much evidence that people often judge the probability of an outcome given a model by the extent to which the outcome represents the model. In addition, Tversky (1977) showed that adding to an outcome (O) a feature (F) that matches a model (M) enhances the match between the outcome and the model. In other words, the match of O&F to M can be greater than the match of O to M. Hence, insofar as people judge the probability of outcomes by their representativeness, being a bank teller and active in the feminist movement would be judged more likely an outcome for Linda than being a bank teller, due to the addition of a feature that is representative of Linda (feminism) to her unrepresentative job. Whereas there is nothing logically wrong with the judgment that being a feminist bank teller is more *representative* of Linda than being a bank teller, judging the conjunctive outcome to be more *probable* than its constituent violates the logically necessary conjunction rule.

Another special case of the extension rule is the disjunction rule, according to which the probability of A-or-B can be smaller than neither the probability of A nor the probability of B, since it contains both. Formally speaking, there is no difference between the three rules (conjunction, disjunction, and extension), because for any pair of events A and B in which B is a subset of A, A can always be represented as a disjunction, one of whose constituents is B, and B can always be represented as a conjunction, one of whose constituents is A. For example, one can argue that the set of bank tellers is a disjunction—of bank tellers who are active feminists with bank tellers who are not. Viewed in this way, Tversky and Kahneman's (1983) results could just as well have been labeled the *disjunction fallacy*. Why then are they regarded as a *conjunction fallacy*? Is this just a matter of arbitrary choice?

Formally speaking, the answer is yes, but psychological considerations favor one view over another. Consider the category *parent* and its subcategory *mother*. One can just as well choose to view *mother* as the conjunction of *parent* with *female* as to

Maya Bar-Hillel and Efrat Neter, Department of Psychology, Hebrew University, Jerusalem, Israel.

Correspondence concerning this article should be addressed to Maya Bar-Hillel, Department of Psychology, Hebrew University, Jerusalem 91905 Israel.

view *parent* as the disjunction of *mother* with *father*. In contrast, the category *bank teller* does not naturally evoke a representation as a union, and certainly not as a union of bank tellers who are active feminists with bank tellers who are not. At the same time, the subcategory of bank tellers who are active feminists can hardly be described except by resort to the conjunction of these two constituents. Indeed, the language does not even contain a single-word label to designate this category. In that sense, the categories *bank teller* and *bank teller and active feminist* are more naturally viewed, respectively, as a unitary category and a conjunction of two categories, than as a disjunction of two categories and a unitary category.

How, then, might one create a category that would be naturally viewed as disjunctive? The simplest possibility to come to mind is to replace the connective *and* used to create conjunctive categories with the connective *or*. This idea must be implemented with caution, however, because the English words *and* and *or* do not always quite correspond to the logical connectives *and* and *or*. First, the English *or* is often understood in its exclusive sense of "A or B but not both," as in "The party will take place next week or the week after." Second, the English *and* can be used to create a union as well as an intersection—the sentences "She invited colleagues *or* relatives" and "She invited colleagues *and* relatives" could be used to describe the same guest list. Third, and most pertinent to present concerns, not all categories that can be meaningfully joined by one of these connectives lend themselves to as meaningful a joining by the other. For example, whereas putting *and* between *bank teller* and *active in the feminist movement* creates a meaningful category, putting *or* between these two category names creates a rather odd one. Similarly, whereas the question, "Is Linda more likely to be a bank teller, or a bank teller and active in the feminist movement?" makes some sense, the question, "Is Linda more likely to be a bank teller, or a bank teller or active in the feminist movement?" sounds to us rather confusing.

Nonetheless, this was precisely the approach taken by all previous attempts to study a disjunction fallacy. In the first attempt to extend the conjunction fallacy to a disjunction fallacy, Morier and Borgida (1984) gave subjects Linda's description and asked them to estimate the probability that (a) Linda is a bank teller, (b) Linda is active in the feminist movement, (c) Linda is a bank teller and is active in the feminist movement, and (d) Linda is a bank teller or is active in the feminist movement.

Wells (1985) took a similar approach. He also gave subjects personality descriptions and asked them for the probabilities of four events: two individual events, their conjunction, and their disjunction. Wells, however, took two precautions to make the resulting disjunction more natural. First, both of the individual events were attitudes toward some national issues, making them of a kind that sounds quite natural when joined by *or*. Second, Wells explicitly added *or both*, to highlight that the *or* was non-exclusive. For example, after describing Jim, subjects were asked for the probability that Jim (a) "favors . . . a U.S. buildup in military strength," (b) "favors the decriminalization of marijuana," (c) "favors [the first] *and* [the second]," and (d) "favors [the first], or [the second], or both." Yet, in spite of these precautions, Wells himself remained concerned "that these sub-

jects misunderstood the union request (e.g., interpreted "or" as a conjunction)" (p. 277).

Biela (1986) took a slightly less formal approach when creating conjunctive and disjunctive categories. After describing to physicians one or two symptoms of an otherwise unknown patient, he asked them to attach degrees of confidence to each of the following diagnoses: (a) "I would predict hypertensive encephalopathy," (b) "I would predict hypertensive retinopathy," (c) "I would predict that either hypertensive encephalopathy or hypertensive retinopathy (or both) is the case," and (d) "I would predict both hypertensive encephalopathy and retinopathy." The subtle variation in formulation between the intended disjunctive (c) and conjunctive (d) diagnosis show Biela's awareness that mechanical joining of the two diagnoses by *or* or *and* is unsatisfactory. Still, (d) might well be understood by a reader to mean "I would give the following degree of confidence to *either* of these diagnoses." In the context of (a), (b), and (c) this possibility becomes less likely, but this comment is intended to show how tricky it is to create categories that are unambiguously disjunctive.

Another difference between categories formed by the connective *or* and categories formed by *and* lies in the difference between matching an instance to A-and-B versus to A-or-B. To be concrete, consider matching Linda to the conjunction "feminist and bank teller," on the one hand, and to the disjunction "feminist or bank teller," on the other. The conjunction requires a single comparison, albeit to a compound event. The disjunction, however, seems to require two comparisons, one for each of the constituent events. Tversky's (1977) theory relating similarity judgments to stimuli's features makes no predictions concerning how two constituent similarity judgments are combined to yield a single similarity judgment for the disjunctive event.

Judging the similarity of the compound event can be sidestepped, however, as it was in Carlson and Yates' (1989) study. Subjects were presented with pairs of events, such as "Syria and Israel will sign a peace treaty by the end of this year" and "The Bill Cosby show will not be one of the top 10 rated TV shows at the end of the season," as well as the conjunction and the disjunction of these events, and they were asked to rank them by their probability of actually occurring within the coming year. But because no model was provided according to which subjects could judge representativeness, the probability of the compound events could only be derived from some kind of combination of the probabilities of its constituent events, not from similarity judgments.

The type of problems used by Carlson and Yates (1989) were termed "probability combination problems" by Gavanski and Roskos-Ewoldsen (1991). These authors recently showed that, when constituent probabilities were controlled for, the rates of the conjunction fallacy were similar in problems in which representativeness could be used to assess the probability of the conjunctive event (e.g., the Linda problem) and in problems in which representativeness could not possibly be used (e.g., probability combination problems). This led them to conclude that the conjunction fallacies "stem primarily from the incorrect rules people use to combine probabilities" and that the "only contribution of representativeness stems from its effects on a

conjunction's component events" (p. 190). Insofar as the disjunctive events used in all studies to date were created by joining constituent events by the connective *or*, the same conclusion could apply to the disjunction fallacy.

The purpose of the present study was twofold: First, to explore the possibility of extending the conjunction fallacy to a more general extension fallacy, while using natural disjunctive categories rather than the somewhat mechanically derived ones used by previous researchers; Second, to explore whether extension fallacies can be obtained even in problems that involve no compound events, hence the fallacy could not possibly result from incorrect combination rules.

Study 1

Method

For this study, we sought categories that would be disjunctive in character, yet defined without the connective *or*.¹ Our solution derived from Rosch's work on categorization (e.g., Rosch, 1978). Recall that the classification of objects into categories is not unique. Thus, the same object can be classified as "artifact," "furniture," "chair," "office chair," or "Eleanor's black leather office chair," to mention but a few possibilities. Yet most people would usually think of it as just "chair." In the hierarchy of nested categories, the level at which objects are commonly classified is known as "the basic level category" (Rosch, 1978). The basic level is not determined arbitrarily. Objects at this level have many features in common—considerably more than those shared by members of higher level categories, but only slightly less than those shared by members of lower level categories. For example, the number of attributes commonly seen as common to *furniture* is 3, to *chair* is 9, and to *kitchen chair* is 10 (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976, Table 2, p. 388).

We capitalized on the observation that basic level categories are typically unitary (i.e., standardly thought of as neither conjunctive nor disjunctive categories), whereas higher levels are often disjunctive (e.g., "furniture" is the union set of "chairs," "tables," "beds," "cupboards," etc.), and lower levels are often conjunctive (e.g., "leather chair"). However, superordinate categories often enjoy a name or label of their own, and are not explicitly defined as compound sets. If a described instance were to match a unitary category more than it matched its superordinate, a disjunction error could potentially arise. Moreover, in the absence of explicit compounding, such error could not be attributed to incorrect combination rules.

Table 1 shows the five types of questions that we devised: (a) The instance is a brief personality description of a student, and the options are fields of university study (Danielle and Oded); (b) The instance is a brief personality description, and the options are places of residence (Gidi and Eldar); (c) The instance is a description of a place, and the options are countries or continents (Gila and Na'ama); (d) The instance is a brief personality description, and the options are social-political movements (Eli and Ze'ev); and (e) The instance is a brief personality description, and the options are causes of death (Yaron and Y.C.).

The target option (i.e., the one designed to be most representative of the instance) was selected to be a basic level category name. The superordinate categories are natural ones, designated by familiar proper names or labels. To strengthen the tendency to view them as a union, their names were prefaced by the words *one of the* or *some kind of*. For example, the basic level cause of death "car accident" was paired with the superordinate "some kind of unnatural cause."

All questions were accompanied by seven options. To give a typical example, the list of fields of university studies provided with the description of Danielle was literature, humanities, physics, natural sciences,

geography or geology, statistics, and political science (not in this order). It is important to note that the lists always included two pairs of a basic level category and its superordinate category (here, the first four options).² Also, there was an option defined with the connective *or* (here, the fifth option). The actual ordering of the various options varied from problem to problem.

Israeli university students taking summer courses (about 60 in advanced economics, 40 in elementary economics, 40 in statistics, 60 in mathematics, 35 in accounting, 15 in sociology) were given from one to four prediction problems, but never two of the same type. On the whole, this subject population is comparable to those identified in other studies (e.g., Tversky & Kahneman, 1983; Wolford, Taylor, & Beck, 1990) as "informed," although subject sophistication was not one of the variables studied here.

There are several concerns a disjunction fallacy might raise. Although some of them were previously raised—and answered—by Tversky and Kahneman (1983) with regard to the conjunction fallacy, in the present context the concerns take on somewhat different twists, and we used somewhat different approaches to deal with them.

First, we were concerned that subjects who see a disjunctive category listed alongside one of its constituents would interpret the disjunctive category in the subtractive sense. For example, given "one of the European countries" alongside "Switzerland," some subjects might implicitly exclude Switzerland from the European countries. Even though Tversky and Kahneman (1983) and Morier and Borgida (1984) ruled out this account of the conjunction fallacy, and it seems unlikely altogether with our natural categories, we thought it prudent to explicitly discourage the possibility.

Tversky and Kahneman (1983) handled this concern by way of what they called "the indirect test" (p. 297). Some subjects were given a list of outcomes that included either the larger category or its subcategory, but not both. When the larger category is not listed alongside one of its subcategories, there is no reason to interpret it as excluding that subcategory. In an indirect test it is impossible to rank "bank teller" before "bank teller and active feminist," because no list includes both options. But across groups, it is possible for "bank teller" to receive a mean rank that is higher than that given to "bank teller and active feminist."

Our approach was simply to include in the subjects' instructions the following explicit cautionary note:

It is quite possible for something to be included in more than one option listed. For example, suppose the list is of foodstuffs, and one option is "frozen foods", while another is "desserts". In this case, do not interpret "frozen foods" as "frozen foods excluding desserts", nor interpret "desserts" as "desserts which are not frozen"—"ice-cream" qualifies as both "frozen food" and "dessert".

The second concern has to do with linguistic conventions. In their discussion of the conjunction fallacy, Tversky and Kahneman brought up the question of how the Gricean concept of cooperativeness (Grice, 1975) applies to a speaker who is less than certain of the truth of his or her assertions. They suggested that such a speaker "may wish to follow the maxim of value: Select the message that has the highest expected value" (Tversky & Kahneman, 1983, p. 312), where "the expected value of a message can be defined as its information value if it is true, weighted by the probability that it is true" (p. 312). With respect to the assertion

¹ Study 1 began as a class project. We gratefully acknowledge the contribution of Anat Ben-Simon, Eyal Gamliel, and Hannan Goldschmidt to the study. We thank Danny Kahneman and Ilan Yaniv for helpful suggestions.

² Actually, there are two exceptions—Hadar ha'Carmel, which is a Haifa neighborhood (Gidi), and North Tel Aviv, which is a Tel Aviv neighborhood (Eldar), but these were not target pairs.

Table 1
Stimuli Used in Experiment 1

Character	Target category	Target superordinate category	Foil category	Foil superordinate category	Description
Danielle	Literature	Humanities	Physics	Natural sciences	Sensitive and introspective. In high school she wrote poetry secretly. Did her military service as a teacher. Though beautiful, she has little social life, since she prefers to spend her time reading quietly at home rather than partying. What does she study?
Oded	Physics	Natural sciences	Literature	Humanities	Did his military service as a combat pilot. Was a brilliant high school student, whose teachers predicted for him an academic career. Independent and original, diligent and honest. His hobbies are shortwave radio and Astronomy. What does he study?
Gidi	Tel Aviv	Dan Metropolitan Area	Hadar ha'Carmel	The North of Israel	23 years old, he wears the latest fashions, and drives a new sportscar. He spends time in discotheques and expensive pubs, and is a social butterfly. He is occasionally mentioned in the gossip columns. Where does he live?
Eldar	Kibbutz Rosh Ha'Nikra	The North of Israel	North Tel Aviv	Dan Metropolitan Area	23 years old, he dresses modestly, and wears sandals year-round, even in winter. He is tanned from outdoor work. He spends his leisure time hiking in the countryside. Where does he live?
Gila	Japan	Asia	Canada	North America	Writes letter home describing a densely populated country, in which modern technology coexists with an ancient local culture. The people, hardworking and competitive, are not inclined to invite strangers into their homes. Where was the letter written?
Na'ama	Switzerland	Europe	Brazil	Latin America	Writes letter home describing a country with snowy wild mountains, clean streets, and flower decked porches. Where was the letter written?
Eli	Peace Now	A peace movement	Gush Emunim	A national movement	39 years old, a professor of Greek Philosophy and Ethics, he holds socialist views. Following the Lebanon War he became politically active, while remaining a "bleeding heart." Where is he active?
Ze'ev	Gush Emunim	A national movement	Peace Now	A peace movement	39 years old, with a red beard and fiery eyes. Married, and a devoted father of five. He teaches Holy Studies in a West Bank settlement. He immigrated from the US 18 years ago. He spends his leisure time hiking in the countryside. Where is he active?
Yaron	Road accident	Unnatural cause	Cancer	Disease	Till he died suddenly at 27, was fully of gait and life. A womanizer, he often drank, and acted wildly in parties and on the road. To his concerned friends he always said, "It couldn't happen to me." What did he die of?
Y.C.	Lung cancer	Cancer	Road accident	Unnatural cause	A woman who smoked over a packet a day for over 10 years. What did she die of?

of disjunctions, it is "misleading (though not false) for a speaker to make a disjunctive assertion where he is in a position to assert one of the disjuncts. Thus, when a speaker asserts A or B he cannot assert A by itself since it is too uncertain" (Adams, 1975, pp. 19-20). For example, I am uncertain when I saw you last, and I assert that "it was either at last year's Psychonomics or when I was in Boston this summer." A listener who assumes Gricean cooperativeness would be justified in inferring from this assertion that I believe the disjunction is fairly likely, but neither of its constituents is sufficiently likely to be asserted by itself.

Applying the value maxim to our problems, a cooperative speaker under uncertainty would only assert, "Danielle is majoring in the humanities" if he or she is not in a position to assert the more informative "Danielle is a literature major." If, on the other hand, one can assert, "Danielle is a literature major," one would not assert, "Danielle is majoring in the humanities," because that would mislead a listener into thinking, counterfactually, that one could not have asserted, "Danielle is a literature major." Hence, the disjunctive assertion, rather than following logically from the more specific assertion, is actually incompatible with it, in the sense that the two would be made under nonoverlapping conditions. The rules of conversational implicature differ from those of logical entailment.

In a betting paradigm (as well as in some legal contexts, most notably contracts), however, logical entailment overrides conversational implicatures. The proposition that "Danielle is majoring in the humanities" is undeniably verified upon learning that "Danielle is a literature major." Even a cooperative speaker, who in everyday conversation would not assert that "Danielle is majoring in the humanities" when "Danielle is a literature major" could be asserted, can legitimately prefer to bet that "Danielle is majoring in the humanities" than that "Danielle is a literature major." The conditions for asserting propositions do not coincide with those for betting on them. In particular, one might sometimes legitimately prefer to assert A than B, yet prefer to bet on B than on A.

To address this concern, we ran two variants of our problems. In the first, subjects were asked to rank the listed categories by the probability that they included the described case as a member. In the second, a different group of subjects ranked the categories by their willingness to bet that the described case was a member in these categories. In this task, no mention whatsoever was made of the word *probability* (thus excluding any confusions this term might engender). In a betting paradigm, the question of why particular outcomes were chosen to be stated and not others is not as relevant as it is for the original task. There may be conditions under which one would be more likely to assert A than A-or-B, but there are no conditions under which a bet on A is more likely to win than a bet on A-or-B.

Results and Discussion

The categories used in Study 1 were not studied in the categorization literature (which typically uses object categories). Hence, a manipulation check was in order, to establish that our target categories really are basic level ones, in the sense that they are at the level where people would tend to locate instances. We asked a separate group of 48 subjects to answer the following questions: (a) What does your best friend study? (100%); (b) Where do your parents live?³ (80%); and (c) To what place in the world would you most like to travel? (75%). As hoped, most respondents answered the questions at the level of the target category, namely they gave answers like (a) physics (rather than natural sciences or quantum theory), (b) Haifa (rather than up north or Hadar ha'Carmel), (c) Japan (rather than the Far East or Tokyo), respectively. Note the shift in location of the basic level that occurs between the questions "Where do your parents

live?" and "Where would you like to travel?" for obvious contextual reasons.

The percentages in parentheses after the questions above indicate how many respondents answered each question at the expected level. In the first question, 11 said, "Does not study at all," so 100% is based on the 37 who responded; in the second question, three people said in general "city" or "in the country," and the rest added a neighborhood name to the town name; in the third question, one person gave a state ("Hawaii"), three people gave city names (e.g., "Moscow"), and the rest gave either just continents (e.g., "Africa") or continents alongside country names (e.g., "South America [Argentina]").

We also checked that our subjects knew that our target categories were contained in their superordinate categories. The same 48 subjects also answered the following questions: (a) In which faculty is the Department of Physics? (96%); (b) In which faculty is the Department of Literature? (98%); and (c) What part of the country is Rosh Ha'Nikra in? (98%). No subjects answered these questions in error, although one or two did not answer at all.

The other categories were not checked (e.g., we did not check that "car accident" is a basic level category whereas "unnatural cause" is its superordinate category, nor that people know that car accidents are unnatural deaths) because the check seemed either superfluous or tricky (see, however, the manipulation check in the following Study 2).

The results of the disjunction study are shown in Table 2. The top row shows the rates of violation of the disjunction rule for the "rank the outcomes by your willingness-to-bet on them" formulation, the middle row shows the rates for the "rank the outcomes by their probability" formulation, and the bottom row combines the two conditions, which of course are normatively equivalent. (The rates in parentheses will be explained later.)

There are small but systematic differences between the probability version and willingness-to-bet version. The rate of the disjunction fallacy is lower under the betting formulation in 8 of the 10 cases (the exceptions being Gidi and Y.C.). However, the decrease averages less than 10% (even the most extreme difference, obtained for Eldar, and the overall difference, fall just short of significance, $z = 1.57$, $z = 1.65$, respectively). The rates of the disjunction fallacy never fall below 50% in either task. We conclude that the disjunction fallacy cannot be accounted for only by conversational implicatures.

The rate of the disjunction fallacy in our problems averaged .64, and never exceeded .83 (Gidi). This is less than the typical rates reported for the conjunction fallacy. Why might this be the case? Perhaps the disjunction rule simply is somehow more compelling than the conjunction rule, although this explanation has a distinct ad hoc flavor.⁴ More likely, note that Tversky

³ We did not ask questions (a) and (b) about the respondents themselves, because all those sitting in a class together at the time of answering the questionnaire studied the same topic and lived in the same city.

⁴ We hasten to note that in both Study 1 and Study 2 we did not select stimuli for presentation on the basis of their effectiveness in eliciting the disjunction fallacy, so the readers are not seeing a selective set of best shots. We did throw out some results, however, as in one problem in

