Cognition 122 (2012) 123–134

Contents lists available at SciVerse ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/COGNIT

On the conjunction fallacy and the meaning of *and*, yet again: A reply to Hertwig, Benz, and Krauss (2008)

Katya Tentori [a,*], Vincenzo Crupi [b,c]

[a] DiSCoF, CIMeC, University of Trento, corso Bettini 31, 38068 Rovereto, Italy
[b] Department of Philosophy, University of Turin, via S. Ottavio 20, 10124 Turin, Italy
[c] Munich Center for Mathematical Philosophy, Ludwig Maximilian University, Ludwigstrasse 31, D-80539 München, Germany

ARTICLE INFO

Article history:
Received 25 July 2010
Revised 9 September 2011
Accepted 14 September 2011
Available online 12 November 2011

Keywords:
Conjunction fallacy
Pragmatic inferences
Probabilistic reasoning

ABSTRACT

In this paper we question the theoretical tenability of Hertwig, Benz, and Krauss's (2008) (HBK) argument that responses commonly taken as manifestations of the conjunction fallacy should be instead considered as reflecting "reasonable pragmatic and semantic inferences" because the meaning of *and* does not always coincide with that of the logical operator $\wedge$. We also question the relevance of the experimental evidence that HBK provide in support of their argument as well as their account of the pertinent literature. Finally, we report two novel experiments in which we employed HBK's procedure to control for the interpretation of *and*. The results obtained overtly contradict HBK's data and claims. We conclude with a discussion on the alleged feebleness of the conjunction fallacy, and suggest directions that future research on this topic might pursue.

1. Introduction

Since the early Eighties, about a hundred scientific papers on the conjunction fallacy (CF) have been published. Such wide interest is easy to understand, as the CF has become a key topic in the fervent debate on human rationality. Indeed, from the very beginning the CF phenomenon has been described as a violation of "the simplest and the most basic qualitative law of probability" (Tversky & Kahneman, 1983, p. 293; but already mentioned in Tversky & Kahneman, 1982, p. 90). The law at issue is the *conjunction rule*, a principle whose compelling nature appears unequivocal when stated formally: $\Pr(p \wedge q) \leqslant \Pr(p)$, i.e., the joint occurrence of a pair of events (p and q) cannot be more probable than the occurrence of anyone of them (e.g., p).

In contrast, what does seem surprising across more than 30 years of research is the recurrence of questions about the validity of CF experiments. A standard line of argument inspired by the pragmatics of communication has been that violation of the conjunction rule need not be irrational if it results from interpreting the experimental task in ways that rob it of normative relevance. The main sources of misinterpretation considered in the literature include participants' understanding of the isolated conjunct p, the term *probable*, and the connective *and*. Many techniques have been developed to control for each of these possible misinterpretations (see Moro, 2009, for a recent review), but none of them has dissipated the effect.

Nonetheless some concerns turned out to be important and should be credited for having fostered improvements in the experimental procedures by which the CF is investigated. To illustrate, the suspicion that the single conjunct p might be interpreted as p-and-not-q (Adler, 1984; Dulany & Hilton, 1991; Messer & Griggs, 1993; Morier & Borgida, 1984; Polizer & Noveck, 1991; but already discussed in Tversky and Kahneman, 1982, 1983) led to more careful control of stimuli, such as explicitly including the statement p-and-not-q in the judgment task along with p and p-and-q. When this technique is applied (as in Tentori, Bonini, & Osherson, 2004; Wedell & Moro, 2008), the rate

* Corresponding author.
E-mail address: katya.tentori@unitn.it (K. Tentori).

0010-0277/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2011.09.002

of violations of the conjunction rule is lower than first reported by Tversky and Kahneman (1982, 1983; the original "Linda" problem) but remains prevalent (e.g., more than 50% for the majority of the scenarios in both studies cited above). Such a pattern makes clear that misunderstanding of the single conjunct cannot be considered the ultimate reason for the occurrence of the CF. However, it also strongly suggests that misunderstanding of the conjunct should indeed be avoided in order to distinguish proper and improper fallacy answers.

In other cases, however, pragmatic factors have been evoked time and again regardless of theoretical remarks and experimental results pointing in the opposite direction. This appears to be the case with the argument advanced by Hertwig, Benz, and Krauss (2008) (hereafter HBK) – previously raised by Ahn and Bailenson (1996), Gigerenzer (1996, 2001, 2005) and Hertwig and Gigerenzer (1999) – that unintended interpretations of the connective *and* may account for (apparent) CF behaviour. As a matter of fact, this concern has been already extensively explored and rejected as unsupported in the literature (see, for example, Bonini, Tentori, & Osherson, 2004; Crandall & Greenfield, 1986; Sides, Osherson, Bonini, & Viale, 2002; Tentori et al., 2004; as well as Moro's review, 2009, specifically devoted to possible sources of misunderstanding in CF tasks).

In what follows, we question the theoretical tenability of HBK's argument as well as the relevance of the new experimental evidence they provide in its support. Subsequently, we reinforce our criticism of HBK on the basis of novel empirical data. Finally, we discuss the alleged feebleness of the CF evoked by HBK, and suggest directions that future research on the CF might pursue.

2. The conjunction rule and HBK's argument

The main point of HBK is that the conjunction rule invokes the logical connective $\wedge$ whereas its experimental test typically relies on natural language conjunctions like English *and*. In contrast to the former, the latter can convey a wide range of relationships between conjuncts (such as *temporal* or *causal* ones) as well as reflect very different set-theoretical operators (such as *union* or *intersection*). HBK argue that depending on which meaning of *and* is assumed, "people may arrive at nearly opposite understandings of a sentence" (p. 741), so that responses commonly taken as manifestations of fallacious reasoning in fact emerge from "reasonable pragmatic inferences" (p. 752). Therefore, HBK conclude, "estimates of the prevalence of genuine conjunction errors in previous studies are quite inflated" (p. 752).

From a theoretical perspective, we find HBK's argument to be affected by two important flaws.

First, the uncontroversial fact (recognized as such in the CF literature ever since Tversky & Kahneman, 1983, p. 302) that the word *and* can have different interpretations across different sentences does not imply anything about its ambiguity within a given sentence. For example, we agree with HBK that *and* in "Mark invited friends and colleagues to his party" is unlikely to be interpreted as involving the *intersection* between the set of friends and the set of colleagues. However, this does not entail that people assign to this sentence multiple contrasting interpretations. Indeed, discussing the very same example, Mellers, Hertwig, and Kahneman (2001, p. 270) pointed out that such *and* "implies a *union*, not an *intersection*", meaning that they do not see any room for equivocation. Should single occurrences of *and* be usually ambiguous, the costs in ordinary conversation would be dramatic, precisely because *and* is "one of the most frequent words in the English language" (HBK, p. 740). As already proved in Tentori et al. (2004), as well as recognized by HBK themselves (pp. 744 and 745), the word *and* in sentences like "Mark has blue eyes and blond hair", for which large CF effects have been observed, is indeed interpreted by virtually all participants in a way that justifies invoking the conjunction rule as a norm.

The second (and major) flaw in HBK's argument is omitting that even when the meaning of *and* is not *exhausted* by $\wedge$, its interpretation often legitimizes application of the conjunction rule all the same. As explained in Tentori et al. (2004), reference to the conjunction rule does not require logical equivalence between *and* and $\wedge$, but only that the interpretation of the *and* statement at issue *implies* the corresponding $\wedge$ statement. A relevant example from Levinson (1983), reported by HBK (p. 747) is: "he turned on the switch and the motor started". Here, the connective *and* may express not only a conjunction between two events but a *temporal* and a *causal* relation that, of course, goes beyond the meaning of $\wedge$. However, if the reader recognizes that both events "he turned on the switch" and "the motor started" must happen for the sentence "he turned on the switch and the motor started" to be true, then the meaning s/he assigns to *and* includes that of the logical operator $\wedge$ in the sense pointed out above. As a consequence, the conjunction rule *can* be properly invoked as a norm.

In this connection, consider Levinson's (1983) treatment of *and*-conjunctions in the pragmatics of language, approvingly referred to by HBK at various places. Levinson noticed that in many cases the interpretation of a *p*-and-*q* sentence will not be limited to the logical conjunction $p \wedge q$, but will add to it a statement of the strongest non-logical (e.g., *temporal*) connection between the conjuncts which is allowed by the interpreter's beliefs about the world. In particular, Levinson's discussion of a conjunction *p*-and-*q* as that reported above involves a series of increasing "informational enrichments" of the purely logical reading of *and*, according to the following pragmatic maxim of interpretation (Levinson, 1983, p. 146; quoted by HBK, p. 747).

> Given *p*-and-*q* try interpreting it as:
> (i) *p*-and-then-*q* [he turned on the switch and then the motor started];
> if successful try:
> (ii) *p*-and-therefore-*q* [he turned on the switch and therefore the motor started];
> if successful try also:

(iii) *p*,-and-*p*-is-the-cause-of-*q* [he turned on the switch, and turning on the switch is the cause for the motor to start].

Despite HBK's claim that these interpretations of *and* are problematic for the assumption of the conjunction rule as a norm, they all clearly imply the bare, logical conjunction $p \wedge q$. Indeed, Levinson's usage of the key concept of informational enrichment leaves no doubt that every subsequent informationally enriched interpretation has any weaker reading as a logical consequence – see, for instance p. 135: "there is an expression e_1, more informative than e_2 (*and thus* e_1 *entails* e_2) [...]" (emphasis added, notation adapted). Accordingly, in Levinson's own account (i) states the plain occurrence of both *p* and *q*, *and moreover* that *q* happens after *p*; (ii) states the plain occurrence of both *p* and *q*, that *q* happens after *p*, *and moreover* that *q* follows (in some sense) from *p*; (iii) states the plain occurrence of both *p* and *q*, that *q* happens after *p*, that *q* follows from *p*, *and moreover* that *q* takes place because of *p*. HBK's resort to Levinson's pragmatics in support of their own line of argument is thus unintelligible. In fact, since all three readings (i)–(iii) imply $p \wedge q$, the conjunction rule unequivocally applies: the probability attributed to *p*-and-*q* by a rational agent must not exceed that attributed to the single conjuncts *p*, *q* (for a step-by-step explanation of this point, see Tentori et al., 2004).

3. The ambiguity between *conjunctive* and *conditional* probability: reconsidering HBK's data

So far, we have pointed out that HBK's theoretical argument is flawed. In fact, in order to vindicate the judgment of *p*-and-*q* as more probable than *p* it is not sufficient to generically evoke a lack of complete overlap between $\wedge$ and the natural language *and*. To consider such a judgment as sound, one has to show that *and* in the specific *p*-and-*q* conjunction under consideration is perceived as ambiguous and that this ambiguity is solved by an interpretation of *p*-and-*q* as not implying *p*.

Let us now consider the empirical evidence that HBK gathered to buttress their argument. HBK present their results as concerning two different CF paradigms originally described by Tversky and Kahneman (1983): the $M \rightarrow A$ paradigm, involving an added hypothesis *A* (*q*, in our current notation) which is more representative than hypothesis *B* (*p*) of a model *M* explicitly supplied or otherwise available; and the $A \rightarrow B$ paradigm, involving an added hypothesis *A* (*q*) which provides a plausible motive to entertain hypothesis *B* (*p*).

We will begin by addressing HBK's treatment of the $A \rightarrow B$ paradigm. (HBK's discussion and empirical data concerning the $M \rightarrow A$ paradigm will be the main topic of the next section.)

In their Experiment 3, HBK found a strong (80%) CF when asking participants to choose the more probable between statements like the following:

p: The percentage of adolescent smokers in Germany decreases at least 15% from current levels by September 1, 2003.

p-and-q: The tobacco tax in Germany is increased by 5 cents per cigarette and the percentage of adolescent smokers in Germany decreases at least 15% from current levels by September 1, 2003.

After the choice task, participants were required to rank-order four possible readings of the statement *p*-and-*q* according to their understanding of this conjunction (specifically, *p*-and-*q* along with the three enrichments of *and*-conjunctions (i)–(iii) mentioned above, i.e., *p*-and-then-*q*, *p*-and-therefore-*q*, and *p*,-and-*p*-is-the-cause-of-*q*). HBK reported that participants whose judgment was consistent with the conjunction rule were much more likely to choose the reading *p*,-and-*q* than participants who departed from the conjunction rule. HBK interpreted this data as evidence that the latter may have departed from the conjunction rule because inferring *causal* and *temporal* relationships and therefore, HBK speculate, judging "a conditional rather than a conjunctive probability" (p. 748).

Of course, had participants actually judged the *conditional* probability of *q* given *p* instead of the corresponding conjunctive probability, then there would have been no fallacy (as already observed by Tversky and Kahneman themselves, 1983, p. 302). However, such a speculation turns out to be unsupported by HBK's own data, for at least two reasons. First and foremost, there is no evidence that participants who violated the conjunction rule evaluated a conditional probability. Actually, 45% of them indicated the purely logical conjunction as the reading which best described their understanding of the *and* at issue. The remaining 55% of participants indicated an enrichment of *and*-conjunctions (i)-(iii). But again, by mere logic, the probability of an enriched *and*-conjunction cannot be higher (actually, it is typically lower) than the probability of the bare logical conjunction $p \wedge q$, because, as pointed out above, the former *implies* the latter. Thus, HBK have simply demonstrated that 45% of their CF responders committed a "standard" CF, while 55% of them committed a violation of the conjunction rule which is even *more remarkable* than usually assumed. Second, the observation that participants who do not violate the conjunction rule were unlikely to perceive strong *temporal* or *causal* relations between the conjuncts *p* and *q* is hardly surprising in light of the literature. The absence of the CF when the two conjuncts are perceived as unrelated or at odds with each other (see, e.g., Tversky & Kahneman, 1983, pp. 305 and 306) suggests that the degree of support (or *inductive confirmation*) between the conjuncts plays a critical role in generating the $A \rightarrow B$ class of CF (Tentori & Crupi, 2009; Tentori, Crupi, & Russo, 2011). As a trivial consequence, people who do not perceive any substantial positive relation between the conjuncts would be less affected by the fallacy.

To prove that participants who chose *p*-and-*q* had judged a conditional rather than conjunctive probability, HBK should have explicitly included something as "*q assuming that p, whether or not p is actually the case*" among the available readings of the statement *p*-and-*q*. A reasoned preference for such a reading from just those

participants who had violated the conjunction rule would have amounted to a direct indication that violations of the conjunction rule may depend on judging a *conditional* rather than a *conjunctive* probability. But, then again, such evidence is not provided by HBK and is unlikely to appear. In fact, the conditional reading hypothesis had been already ruled out by previous CF studies (e.g., Bonini et al., 2004) and is incongruent with the fact that none of HBK's participants has been reported as complaining that a genuine "conditional reading" of the *and* at issue was not offered as an option.

HBK claim that a direct test of the hypothesis of violations of the conjunction rule may depend on judging a conditional rather than a conjunctive probability comes from their Experiment 4. In this experiment, they presented two groups of participants with matched conjunctive (p-and-q) versus "conditional" (q,-assuming-p) statements, respectively. The probability estimates provided by the two groups did not significantly differ, leading HBK to conclude (p. 749) that "this is consistent with the hypothesis that (some) respondents in [HBK's] Experiments 3 and 4 and in Sides et al.'s (2002) experiments have interpreted the and-conjunctions in terms of cause-effect statements, and as a consequence judged conditional rather than conjoint probabilities." The main problem with this conclusion is that results of HBK's Experiment 4 cannot be directly related to any violation of the conjunction rule since in that experiment no CF task was involved.[1] Moreover, many different explanations are plausible for the reported lack of a difference between conjunctive versus conditional probability estimates. These explanations are ignored by HBK and include, among others, a mere overestimation of the conjunctive probability (as it would be in a proper CF) or an underestimation of the conditional probability. As to the latter point, it is also worthwhile to notice that the accuracy and provenance of conditional probability estimates are themselves somewhat controversial in the literature (see, for example, Koehler, 1996; Villejoubert & Mandel, 2002; Zhao, Shah, & Osherson, 2009; but also Hertwig & Chase, 1998, p. 324).

4. The ambiguity between *conjunction* and *disjunction*: A new empirical test

In the previous section, we raised concerns as to HBK's results allegedly suggesting that violations of the conjunction rule may depend on judging a *conditional* rather than a conjunctive probability. Indeed, what HBK's Experiment 3 proved is only a high rate of violations of the conjunction rule by participants who interpreted p-and-q as implying the logical conjunction $p \wedge q$, thus a genuine, plain and large CF. In the present section, we consider another set of results reported by HBK in favor of the same general thesis that the CF stems from a misinterpretation of the connective *and*.

In the case of the $M \rightarrow A$ paradigm (Tversky & Kahneman, 1983), HBK ascribe departures from the conjunction rule to a conflation between the probability of a *conjunction* and that of a *disjunction*. This hypothesis of HBK, just like the one previously discussed, has already been considered and ruled out by some studies. Tentori et al. (2004), for example, observed a large amount of violations of the conjunction rule by participants who overtly endorsed the implication from p-and-q to p, i.e., participants for which the conjunction/disjunction conflation is simply out of the question. Yet HBK presented new data in support of their hypothesis, therefore we will consider it more specifically on this novel ground.

In their Experiment 1, HBK presented 119 participants with a version of the Linda problem, and randomly assigned them to three different groups who were asked to estimate the number of "bank tellers" (group 1), "bank tellers and active feminists" (group 2), and "feminist bank tellers" (group 3), respectively. Subsequently, all participants were asked to shade the area corresponding to the quantity they had just estimated on a Venn diagram, of which HBK report (p. 743) that "consisted of two partly overlapping circles (one of which was labeled 'bank tellers' the other 'active feminists')." HBK found that the mean estimated number of "bank tellers and active feminists" provided by group 2 (21.8) was significantly higher than that of "bank tellers" provided by group 1 (7.1), thus replicating the between-subjects conjunction effect already described in previous studies (see, for example, Tversky & Kahneman 1982, 1983).[2] However, HBK also found that, in group 2, the mean estimate provided by those participants (24%) who shaded the *union* of sets in the subsequent task was significantly greater than the mean estimate provided by the participants (76%) who shaded the *intersection* of sets (59.1 versus 10.2, respectively). The former but not the latter estimate differs significantly from that of "bank tellers" provided by group 1, suggesting to HBK that the conjunction effect observed by comparing groups 1 and 2 was entirely due to those participants in group 2 who had interpreted *and* as conveying a *union* of sets. HBK claim further support for their analysis from the observation that "about half of respondents in group 2 (27 out of 59) explicitly asked the experimenter – during the estimation task – how the *and* was meant to be understood" (HBK, p. 743). HBK concluded that the ambiguity of the conjunction in "bank tellers and active feminists" explains the higher mean estimate given by group 2 to this conjunction with respect to that given by group 1 to the conjunct "bank tellers", "causing the appearance of a conjunction effect" (HBK, p. 743).

The idea of controlling how the conjunction is interpreted *after* performance on a CF task is not new in the literature, and we agree with HBK that it may provide important hints on the nature of the CF phenomenon. Of course, the way this control is carried out is crucial for

[1] Also note that, in Experiment 4, HBK rely on proving a negative (i.e., the absence of a difference). Finally, an extensive overlap between conditional and conjunctive estimates is itself problematic for HBK, for surely on their account not *all* participants in the conjunctive-statement group should be giving the conjunction a conditional reading. For example, extrapolating from HBK's interpretation of the results from their Experiment 3, only about a half of the participants should have judged a conditional probability when asked for a conjunctive one.

[2] When a between-subjects design is employed, the CF is sometimes referred to as "conjunction effect" to underline that non-normative responses are documented from a comparison between groups (e.g., in Tversky & Kahneman, 1983).

the reliability of its results. In particular, to provide insight into earlier responses, it should be as invulnerable as possible to observer's expectancies, as well as to participants' inferences on the aim of presenting the control task. To fulfill these constraints, Tentori et al. (2004, pp. 471–472 and 474–475) employed written *implication questions* within a task that was not easily associable to the earlier CF task. On the contrary, the shading task employed by HBK is quite transparent in its goal and can be quite directly related to the earlier CF task, thus allowing for possible carryover effects from the former to the latter. For example, participants in HBK's group 2 who provided a high estimate in the CF task could have been reluctant to shade the intersection because it appeared as a relatively small geometrical area, thus somewhat inconsistent with their high estimate. HBK report that a similar concern was raised by one of their reviewers and reply as follows: "Not having estimated the single event and the conjunction, no mistake had occurred. Consequently, it is not possible that participants' shading of the Venn diagram was an attempt to rectify their earlier judgment." (p. 743, footnote 3). We do not find HBK's point convincing, for even if their CF task was between subjects (and therefore a single participant could not be right or wrong), participants who had provided a high estimate could still have felt uneasy in shading a small area (intersection of sets). Therefore, the risk of a carryover effect in the shading task needs to be considered. A comparison between the answers provided in the same shading task by participants who had versus had not previously estimated the corresponding conjunction could be a way to control for (and possibly rule out) such a potential carryover effect. Alas, HBK asked only their group 2 to shade the area corresponding to "bank tellers and active feminists", while their group 1 had to shade the area corresponding to "bank tellers". This choice makes it impossible to rely on group 1 as a control for group 2 concerning the possible carryover effect at issue.

A concurrent source of confound in the shading task could be the size of the overlap between circles in the Venn diagram. Participants might have been discouraged from shading it just because they happened to perceive the overlap as implausibly small. This possible confound is not independent from the one discussed above: the less the two circles overlapped, the stronger was the risk of a carryover effect for participants who had assigned a high estimate to the conjunction in the estimation task. Unfortunately, HBK did not describe the graphical details of the Venn diagram they presented, nor the specific instructions they used to introduce the shading task to their participants.

A final concern regarding HBK's procedure arises from the report (p. 743) that half of their participants "explicitly asked the experimenter – during the estimation task – how 'and' was meant to be understood". Participants in all our data collections with CF tasks (Bonini et al., 2004; Tentori & Crupi, 2009; Tentori et al., 2004, 2011), although interviewed either individually or in very small groups (2–3 participants at maximum), never asked such a clarifying question, nor is it ordinarily reported in papers describing CF experiments with the Linda scenario (there is no mention of it, for example, in Tversky & Kahneman, 1982, 1983; Jones, Taylor Jones, & Frisch, 1995; Hertwig & Chase, 1998; Mellers et al., 2001). Such a high incidence (50%) of requests for clarification in the very context of an experiment aimed at testing the ambiguity of *and* is quite striking and suggests that participants may have felt somewhat under pressure.

4.1. Experiment I

To control for the possible confounding factors mentioned above as possibly affecting HBK's procedure, we carried out our Experiment I.

4.1.1. Method

One hundred and twenty students from University of Ferrara (Italy) participated in this experiment (average age: 24 years, 58 females). They were enrolled in different degree programs, including Medicine, Dentistry, Pharmacy, Biology, Geology, Chemistry, Engineering, Economics, Maths, Italian literature, Arts, and Architecture. Participants were not paid for their participation. They were interviewed individually and randomly divided into two groups.

As in HBK, the two groups were introduced to the classical Linda scenario in a frequency format (see Appendix A for a detailed description of the stimuli used in Experiment I) and asked to estimate the number of "bank tellers" (group 1) and "bank tellers and feminists" (group 2), respectively. However, unlike HBK, we then asked all participants to shade in a Venn diagram the area corresponding to the conjunctive statement "bank tellers and feminists".[3] As pointed out above, participants from group 2 in HBK's experiment who had provided a high estimate in the CF task could have been more reluctant to shade the intersection for consistency with that estimate. By asking also group 1 to shade the area corresponding to "bank tellers and feminists", unbiased interpretations of *and* as a disjunction (with equal rate expected across both groups) can be effectively disentangled from carryover effects of the earlier estimating task (which should be obviously absent in group 1).

The Venn diagram comprises two partly overlapping circles, one labeled "bank tellers", the other labeled "feminists" (see Appendix A for the exact dimension and overlap of these two circles as well as the exact wording). To make the diagram as clear as possible, we also used different colors for the borders as well as for the corresponding labels.

Finally, to avoid any experimenter-expectancy effect, the data were collected by a colleague working in another field (neuroimaging of motor control) who was blind to the research hypotheses.

5. Results

The results of Experiment I are reported in Table 1.

[3] We still consider the use of written implication questions (Tentori et al., 2004) a much more reliable control for the interpretation of *and*. To allow comparability with HBK's results, however, we adopted the shading task in the present experiment.

Table 1
Results of the Experiment I (p = "bank tellers"; q = "feminists"). For both the groups 1 and 2, the number (and percentage) of participants who shaded each specific area in the Venn diagram and the corresponding average estimates in the CF task are reported. The exact estimates of the six participants who did not shade the intersection of sets are given in square brackets.

Shaded area	Group 1		Group 2	
	N	Average estimate (p)	*N*	Average estimate (p and q)
$p \wedge q$	59 (98%)	19.20	55 (92%)	36.04
$p \vee q$	–	–	2 (3%)	75 [50; 100]
subset of $p \wedge q$	1 (2%)	20 [20]	2 (3%)	10 [0; 20]
$p \wedge \neg q$	–	–	1 (2%)	30 [30]
	60	**19.22**	60	**36.37**

The average estimate of "bank tellers" provided by group 1 was 19.22 while the average estimate of "bank tellers and feminists" in group 2 was 36.37. The difference between these two estimates is statistically significant ($t(118) = -3.5$, $p < .001$), confirming a between-subjects effect compatible with a violation of the conjunction rule. To qualify this difference as a manifestation of non-normative reasoning versus reasonable pragmatic inferences arising – as HBK claim – from an understanding of *and* that is different from conjunction, we need to understand how *and* was interpreted by participants. The area in the Venn diagram corresponding to the intersection of the "bank tellers" and "feminists" sets was shaded by 98% of participants in group 1 and by 92% of participants of group 2. Therefore, almost all participants in group 1 and a large majority of participants in group 2 interpreted *and* as a conjunction.

The picture is unaffected by considering only those participants who shaded the intersection of sets. The average probability estimates for the two groups are 19.20 and 36.04, respectively, and again the difference between them is highly significant ($t(112) = -3.4$, $p = .001$).

These data strongly support the hypothesis of a genuine departure from the conjunction rule, leaving no room for HBK's conclusion (p. 743) about the CF effect being only "apparent" because of the ambiguity of *and*.

Moreover, among the six participants (out of 120) who did not shade the intersection of sets, only 2 indicated the union of sets as their interpretation of *and*. They were both in group 2 (i.e., the group first asked for an estimation of the conjunction), and their responses were 50 and 100. At least for the latter participant, a carryover effect from his high estimate to the subsequent shading task is plausible. In any event, the fact that no participant in group 1 selected the union of sets clearly indicates that, when these kinds of carryover effects are eliminated within a proper control group, the conjunction at issue proves not to be ambiguous in any relevant sense.

5.1. Experiment II

As a test of robustness and generalizability, we ran a second experiment using a different scenario and a different CF task.

5.1.1. Method

Sixty students (average age: 22.7 years, 39 females) from the same population of Experiment I participated in Experiment II. None had participated in Experiment I. Participants were not paid and were interviewed individually. Once again, the data were collected by an experimenter who was blind to the research hypotheses (the same as for Experiment I).

Table 2
Results of the Experiment II (p = "blond hair"; q = "blue eyes"). The number (and percentage) of participants who shaded each specific area in the Venn diagram and the corresponding choice in the CF task are reported.

Shaded area		Choice		
	N	p	p and q	p and $\neg q$
$p \wedge q$	59 (98%)	23 (38%)	35 (58%)	1 (2%)
$p \vee q$	–	–	–	–
Subset of $p \wedge q$	1 (2%)	–	–	1 (2%)
	60	23	35	2

Participants were presented with a choice version of the Scandinavia problem (Tentori et al., 2004) in frequency format, followed by a shading task (see Appendix B for a detailed description of the stimuli used in Experiment II). A *choice* task allows each participant to compare directly the conjunction (p and q) and the single conjunct (p). Violation of the conjunction rule in such a within-subjects design is what Tversky and Kahneman (1983, p. 298) properly define a CF.

6. Results

As displayed in Table 2, the majority of participants (58%) chose the conjunction "blond hair and blue eyes" as the most probable option. All of them indicated the intersection of sets in the subsequent shading task, discrediting any reading of their choice as anything but a proper CF. Moreover, if we consider the totality of the participants, none shaded the union of sets, indicating that *and* was not perceived as a disjunction.

These results are in line with those of Experiment I and strongly support the hypothesis of a genuine violation of the conjunction rule, documenting one more time that there are clear cases of CF that cannot be ascribed to participants misinterpreting *and* as actually expressing a disjunction or any other relevant meaning.[4]

It is also worth noting that none of the participants in both experiments described above pointed to the conjunctions at issue as ambiguous nor asked how *and* should have been interpreted. This is coherent with

[4] Indeed, the same result was already provided in Tentori et al. (2004) (see more on this point later).

Tversky and Kahneman's (1982, p. 95) observation that, when debriefed, subjects who had committed conjunction errors generally agreed that their judgments were mistaken.

We will not venture to provide a specific explanation of why the results of the shading task in both our experiments do not replicate HBK's results. Especially as HBK do not provide exact information about how they collected their data nor on the characteristics of the Venn diagram they used. Among the possible interpretations, a concern to be taken seriously is a carry-over effect induced by the participants' feeling of a pressure to appear consistent in their judgment. If confirmed, it would restate that control for pragmatic concerns should extend beyond words involved in the stimuli.

7. A feeble fallacy?

HBK lamented (p. 745) a "feebleness of violations of the conjunction rule that appear to be a function of myriad variables including response format (probability versus frequency), response mode (ranks versus estimates), and the presence or absence of filler items [...] and of semantic and pragmatic ambiguities."

In this statement, HBK lump together factors that have hugely different implications for the CF debate. By the consideration of alleged "semantic and pragmatic ambiguities" the *reality* of the CF is disputed: if CF results rested on semantic and pragmatic ambiguities, they would not have documented a real reasoning fallacy. On the other hand, variables such as the specific probability task employed (the "response mode" in HBK's terminology) or the stimuli structure (e.g., the "number of fillers") may at best affect the *generalizability* of the CF: the more the CF is sensitive to these variables, the less it is generalizable.

Of course, both these issues are interesting, yet they should not to be conflated. To begin with, a discussion on generalizability of the CF is expected to presuppose recognition of its reality, for it clearly makes no sense to complain about the low generalizability of a phenomenon while arguing that it does not actually exist. Still, this seems to be the case with HBK. Apparently, they both reject the CF as a real fallacy and at the same time question its generalizability. Second and more important, any conclusion about the reality of the CF drawn from results concerning its generalizability reflects a severe misunderstanding of what a fallacy – and, more generally, a cognitive phenomenon – is. To gain the status of a real fallacy, a reasoning bias need not always have the same strength. As already clarified by Kahneman and Tversky (1996), it is in the nature of both visual and cognitive illusions that conditions exist under which the correct answer is made more transparent. These conditions may affect participants' general cognitive load or provide cues that make the comparison between hypotheses easier. For example, CF rates have been sometimes reported to be lower when an estimation task is employed in absence of fillers (see more on this later on, though). But this does not imply in any way that the strong CF observed under different conditions – such as, for example, choice or ranking tasks – is not real. As relevant as it can be, the import of modulating factors concerns the generalizability issue only, and should be treated accordingly.

Let us illustrate this crucial logical and theoretical point by a further remark. Tversky and Kahneman (1982, p. 305) reported no CF when Linda's personality sketch was omitted from the otherwise usual scenario. Apparently, thus, the relationship between the background information provided and the conjuncts plays a key role in determining CF. (For a detailed discussion of these factors see Crupi, Fitelson, & Tentori, 2008; Tentori et al., 2011.) Although extremely interesting, this clearly has no implication for the reality of the CF, unless one is willing to claim that, by omitting the background scenario, the conjunction between the very same conjuncts ("bank teller" and "active feminist") suddenly becomes unambiguous. To our mind, such an attempt would be as hopeless as the search for a cognitive illusion that is immune to any possible modulation.

An example of argumentation in which HBK mix reality and generalizability concerns is summarized in the following. At p. 744, HBK state that the stimuli in Tentori et al. (2004) did not include ambiguous conjunctions: "Tentori et al. (2004) also made the case that a majority of their participants understood that *p-and-q* implies *p* but nevertheless attached greater probability to the former compared with the latter" (notation adapted). Soon afterwards, however, HBK reject Tentori et al.'s (2004) conclusion that these experiments provide evidence for the CF when the ambiguity of the conjunction is out of the question, saying that "one possible alternative [conclusion] is related to [the fact that] respondents were not instructed to estimate the groups' frequencies but to select the most numerous one" (p. 744). How HBK relate the task employed (a generalizability issue) to the ambiguity of *and* (a reality issue) is not clear to us, but it seems to imply that in their opinion the choice between hypotheses is not a proper probability task, and a "genuine" fallacy in reasoning about chance can only be demonstrated by an estimation task. In fact, later on (p. 744) they speculate that ranking or choice tasks as those in Tentori et al. (2004) "may have triggered participants to assume that some of the groups are mutually exclusive", suggesting still another pragmatic concern, i.e., the *conversational implicature* by which *p* could be read as *p*-and-not-*q* (see the second paragraph in Section 1). Yet soon afterwards HBK acknowledge that Tentori et al. (2004) "discourage such readings" by inserting the statement *p*-and-not-*q* as an available choice option (p. 744). So, once back to the reality issue, they recognize again that the stimuli in Tentori et al. (2004) were not ambiguous, despite the choice task. Do HBK finally conclude that the results of those experiments represent evidence for a genuine CF? Actually, they do not. Instead, they simply turn to a new paragraph introducing their Experiment 2, in which they employ the conjunction presented in Tentori et al. (2004), and precisely because it "avoids the ambiguity of whether it denotes the union (addition) or the intersection of both sets." (p. 745). HBK mention once again Tentori et al. (2004) at the end of their paper (p. 751), where they restate their strong disagreement with the conclusion that Tentori et al.'s (2004) results point to a genuine error in reasoning about chance. Here, HBK's motivation is as follows: "Experiments 3 and 4 provide evidence that suggests that at least some people in Tentori et al.'s (2004) tobacco-tax

task have inferred – in the process of comprehending the sentence – a causal meaning, and, therefore, appear to have estimated the conditional rather than the conjunctive probability." Notice that in Tentori et al.'s (2004) experiments there is neither a tobacco-tax scenario nor, as HBK themselves remark just a few pages ahead, any plausible source of semantic or pragmatic ambiguity.

Such a persistent skepticism on the reality of the CF when possible misunderstandings about the conjunction are considered to be ruled out by all parties appears to us unmotivated. This is even more remarkable as convergent results in favor of the reality of the CF were also reported in other studies, e.g., in Bonini et al. (2004) with a betting task.[5] Here we employed HBK's own procedure to control for the possible ambiguity of conjunctions, and replicated Tentori et al.'s (2004) results one more time. Maybe now it will be eventually agreed that the data do point to a genuine (and elementary) error in reasoning about chance. We certainly submit that this conclusion is very strongly supported.

While HBK's criticism concerning the reality of the CF is unconvincing in light of a large amount of evidence available, one might still wonder if HBK's statements about the low generalizability of the CF are better grounded on existing studies. The issue is not central to our concerns in this paper, so we won't go into much detail. Suffice it to mention at least some of the many results suggesting a conclusion that is quite different from HBK's.

The CF has been documented not only in a variety of *hypothetical scenarios* (as those considered so far), but also in many *real-life domains*, including diagnosis and prognosis in clinical settings (e.g., Garb, 2006; Rao, 2009; Tversky & Kahneman, 1983), professional auditing (e.g., Frederick & Libby, 1986; Ho & Keller, 1994), forecasts of sports results (e.g., Nilsson & Andersson, 2010; Teigen, Martinussen, & Lund, 1996a), effects of government policies (e.g., Bonini et al., 2004), and political outcomes (e.g., Lee, Grothe, & Steyvers, 2009; Teigen, Martinussen, & Lund, 1996b).

The CF has been documented with various kinds of *participants*: not only thousands of university students enrolled in the most disparate programs, but also children (e.g., Agnoli, 1991; Davidson, 1995; Fisk & Slattery, 2005), laypeople (e.g., Hertwig & Chase, 1998), experts (e.g., Adam & Reyna, 2005; Frederick & Libby, 1986; Garb, 2006; Ho & Keller, 1994; Tversky & Kahneman, 1983), and statistically sophisticated individuals (e.g., Tversky & Kahneman, 1982, 1983).

The CF has been documented with different *tasks*, such as choice (e.g., Tentori et al., 2004; Wedell & Moro, 2008), ranking (e.g., Sloman, Over, Slovak, & Stibel, 2003; Tversky & Kahneman, 1983), money allocation (e.g., Bonini et al., 2004), and different estimation procedures (e.g., Nilsson & Andersson, 2010; Nilsson, Winman, Juslin, & Hansson, 2009; Tversky & Kahneman, 1983; Weddel & Moro, 2008). The first three tasks are within subjects by their very nature. Estimation, on the other hand, can be used in both between- and within-subjects designs. In the former case (e.g., HBK's Experiment 1), results largely match those reported with the other tasks. When estimation is employed in within-subjects designs, CF rates have been sometimes reported to be broadly in line with other tasks (e.g., the CF rate is 65% in the *heart attack* scenario from Tversky & Kahneman, 1983, p. 308; also see Nilsson et al., 2009), and sometimes lower, yet still above zero (e.g., ranging from 33% to 55% in Sloman et al. (2003); with an overall level of more than 40% in Wedell and Moro (2008); and 42% in study 1 of Hertwig and Chase (1998)).

The CF has been documented with different *stimuli structure*, for example with (e.g., Mellers et al., 2001; Sloman et al., 2003) as well as without fillers (e.g., Tentori et al., 2004: Wedell & Moro, 2008).

The CF has been documented as being resistant to the introduction of dynamic *feedback* as well as monetary *incentives* (Zizzo, Stolarz-Fantino, Wen, & Fantino, 2000).

The CF has been documented to be also resistant to *linguistic training* aimed at providing a definition of "conjunction" as well as improving participants' accuracy in distinguishing proper conjunctions from other meanings that may be conveyed by the word *and* (Crandall & Greenfield, 1986).

The CF has been documented with problems framed in both *probability* and *frequency formats*. In particular, although the CF has been claimed to be "reduced, or even made to disappear, when participants are given frequency information and asked for frequency judgments" (e.g., Hertwig & Gigerenzer, 1999, p. 299), several studies exist in which indistinguishable CF rates have been observed with stimuli framed in terms of probabilities versus frequencies (e.g., Jones et al., 1995; Sloman et al., 2003; Tentori et al., 2004; Weddell & Moro, 2008).

Given such an impressive consistency and convergence of results, it is no surprise that the CF has been repeatedly claimed to be *robust* and *easy to replicate* (e.g., Gavanski & Roskos-Ewoldsten, 1991; Teigen et al., 1996a, 1996b; Epstein, Donovan, & Denes-Raj, 1999; Sides et al., 2002; Stolarz-Fantino, Fantino, Zizzo, & Wen, 2003; Nilsson, 2008; Wedell & Moro, 2008; Lee et al., 2009).

8. Concluding remarks

In this paper, we have considered the theoretical arguments and empirical results that HBK provide in favor of their interpretation of the CF as a manifestation of reasonable pragmatic and semantic inferences prompted by a misunderstanding of the *and* conjunction.

On the theoretical side, HBK's argument has two important flaws: first, a conflation between the multiple meanings of *and* across different sentences and the alleged ambiguity of *and* within one given sentence; and second, apparent unawareness of a plain logical point, i.e., that meanings of *and* which are strictly stronger or "richer" than $\wedge$ (such as those involving a temporal or causal sequence) do imply the conjunction rule of probability theory as a normative constraint on rational judgment all the same.

[5] The betting task has the merit of avoiding locutions involving the word "probability" whose meaning has been hypothesized to be another source of misunderstanding in CF experiments (see for example, Hertwig & Gigerenzer, 1999). In particular, participants in Bonini et al. (2004) preferred to bet on a pair of events (p-and-q) even if they could bet at the same stakes on a single event (p) included in that pair. The same answers were obtained even when the sentence "both events must happen for you to win the money placed on this bet" was added to each pair of events as to eliminate possible misunderstandings of the conjunction at issue.

This latter major theoretical flaw reverberates on HBK's reading and use of their data concerning the $A \rightarrow B$ paradigm. In fact, no compelling evidence is provided that fallacious judgments in this paradigm can be explained away by postulating conditional rather than conjunctive probabilities having been judged. As argued above, HBK's Experiment 3 proves the opposite, if anything, while HBK's Experiment 4, apart from a number of methodological problems (see footnote 4 above), is simply inconsequential as it does not address any CF task.

Concerning the $M \rightarrow A$ paradigm, we carried out two novel experiments, adopting HBK's Venn diagram shading control procedure in a between-subjects estimation task (Experiment I) and in a choice task (Experiment II). In both cases, to counter possible experimenter-expectancy effects, data collection was carried out by an experimenter blind to the research hypotheses. All the CF responses were provided by participants who interpreted the *and* as a conjunction based on results from the shading task. Moreover, the allegedly critical interpretation of *and* as a disjunction only occurred in a tiny minority of cases (Experiment I), if ever (Experiment II). These results precisely replicate those obtained in previous studies (as Bonini et al., 2004; Tentori et al., 2004), in which possible misunderstandings of the conjunction were ruled out by means of other control procedures, and show one more time that there are clear cases of CF that cannot be ascribed to the ambiguity of *and*.

We conclude that the ability of the human mind for subtle pragmatic and semantic inferences concerning the interpretation of *and* – as striking and interesting as it is – simply does not undermine the CF as a real, elementary and prevalent bias in human judgment under uncertainty. Of course, this is not to say that all experiments carried out on the CF in the last 30 years were completely immune to semantic or pragmatic ambiguities. Some of them – as in principle any psychological experiment involving human participants – might well have been affected by these confounds. Indeed, we firmly believe that pragmatic concerns played a critical role in helping experimenters to refine and improve experimental techniques and stimuli over the years. Anyway, what is relevant for the present debate is the following: does the CF disappear once potential semantic and pragmatic ambiguities are controlled for? Results from a number of controlled studies – including Tentori et al. (2004), Bonini et al. (2004) as well as the experiments reported here – strongly imply a negative answer.

In reviewing this paper, Dr. Hertwig provided the following comment on the failure to replicate HBK's results in the shading task:

> We [HBK] asked our participants to "shade that part of the diagram that represents – in your understanding – the quantity 'bank tellers and feminists'." In contrast, Tentori and Crupi used the instruction: "Shade with the pen the area corresponding to women who are 'bank tellers and feminists'." The key differences between these instructions are that Tentori and Crupi's (i) makes no reference to the estimate that people had just given and (ii) does not mention that people are supposed to express through their shading their understanding of this quantity. I can't say I am surprised that people are able to shade the area corresponding to women who are bank tellers and feminists when asked to do so!

We find this argument unconvincing. To begin with, consider point (i). HBK's shading instructions (reported in Dr. Hertwig's quote) – exactly as ours – make no explicit reference to the estimate that participants had just given, therefore this cannot account for the absence of disjunctive interpretations of *and* in our experiments as contrasted with HBK's. Moreover, even if in the shading task there is no explicit reference to the CF task, we followed HBK in presenting the two tasks one just after the other while referring them to the same conjunction, therefore it is quite implausible that participants perceived those tasks as "unrelated". Finally, notice that for a control task to be reliable, participants need not be aware of its purpose. Quite the opposite, the transparency of the shading task is exactly what makes it vulnerable to possible carryover effects (for this reason, as indicated in footnote 3, we adopted the shading task only for the sake of comparison with HBK's results, but we still consider the implications questions used in Tentori et al. (2004) a better control procedure). In his point (ii), Dr. Hertwig's suggests that the presence of the words "in your understanding" would have shifted participants' responses from the intersection to union of sets, thus revealing the alleged ambiguity of the conjunctive statement. We doubt that such a shift would obtain, as nothing in our instructions encourages participants to perform the shading task in a perspective other than their own. However, in case the presence of the words "in your understanding" had some specific impact, a symmetric concern could well be raised. As already pointed out above, the shading task was quite transparent, it came just after the CF task, and referred to the same conjunction. Thus, one might just as well speculate that the words "in your understanding" in HBK's shading task instructions have been interpreted by participants as an invitation to be consistent with the estimate they had just provided in the CF task, thus fostering a carryover effect. If so, HBK's very results in the shading task could have been induced by a pragmatic misunderstanding between experimenter and participants.

Anyway, it is not our goal to go any further in trying to understand why we did not replicated HBK's results in the shading task. This is because the shading task is just a tool, and not the focus of the present debate. More relevant to the reality of the CF is that there are conjunctions (as the statement "blond hair and blue eyes" employed in Tentori et al. (2004), or in the Experiment II of this article) which are considered unambiguous by all parties (see HBK, p. 745), but still judged more probable than the corresponding single conjuncts by the majority of people. This result is wholly inconsistent with the hypothesis that the CF is just an artifact resulting from a misinterpretation of the conjunction and should not be clouded by pragmatic concerns that apply to other stimuli or considerations concerning the generalizability of the CF to different kinds of tasks. A never-ending revival of old contentions on the reality of the CF despite contrary evidence serves no scientific purpose. We urge, therefore, that the concern about CF

being just an artifact of some misunderstanding between experimenter and participants about the conjunction be clearly recognized for what it actually is: A fascinating hypothesis which has been extensively explored and repeatedly disconfirmed.

So what are some fruitful themes for future research on the CF? Ever since the earlier relevant investigations, probabilistic fallacies have been of interest as providing valuable hints to a comprehensive view of human reasoning. In consideration of this original aim and also of the remarkable amount of research that the CF has prompted, what is surprising is the lack of a generally accepted explanation of the phenomenon, as pointed out by several observers in recent times (e.g., Fisk, 2004; Jarvstad & Hahn, 2009; Nilsson et al., 2009). Therefore, in our view, the main extant topic concerning the CF is an account of the phenomenon displaying adequate explanatory and predictive power. In particular, we do agree with some critics that interpretations based on informally and fuzzily characterized heuristics remain of limited value in that they are resistant to clear-cut empirical test (see, e.g., Birnbaum, Anderson, & Hynan, 1990; Gigerenzer, 1996). Accordingly, more and more effort is being placed in turning from heuristics to sharper models that clarify the antecedent conditions of the reasoning patterns observed (for different attempts in this direction see Kahneman & Frederick, 2002; Lagnado & Shanks, 2002, and Shafir, Smith, & Osherson, 1990). Among the most intriguing alternatives to the traditional heuristic approach, there are explanations which ascribe the CF to (non-normative) averaging rules as applied to the probabilities of the conjunction's constituents (e.g., Nilsson et al., 2009) or to random error in the judgment process (e.g., Costello, 2009). While very different from each other, these proposals all imply that CF rates would rise as the perceived probability of the added conjunct does. A different reading of the phenomenon has also been advanced, based on the notion of inductive confirmation as defined by contemporary Bayesian theorists (Crupi et al., 2008). This predicts that the CF depends on the added conjunct being perceived as inductively confirmed rather than highly probable. Should such a dependence of the CF on intuitive assessments of confirmation be proved, it would shed new light on the subtle connections between probabilistic fallacies and inductive reasoning in human cognition. Future experimental investigations might thus neatly dissociate the relevant variables and predictions of extant competing accounts of the CF. Needless to say, stimuli will need to be devised that appropriately control for possible pragmatic and semantic misunderstandings in their wording, while not being unduly vulnerable to other sources of confound either. Hopefully, on these conditions, real progress might lie ahead in explaining this fallacy.

Acknowledgements

This research was supported by a Marie Curie grant for the project INDUCT (221350), by a MIUR Grant "Inductive reasoning: Models, behavior, and applications" (PRIN, No. 20083NAH2L), and by the Alexander von Humboldt Foundation.

Appendix A

Stimuli used in Experiment I

Page 1

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.
Imagine there are 100 women like Linda.
How many of them do you think are bank tellers [bank tellers and feminists]? _____ out of 100.

Page 2

In the picture below there are represented respectively:

- The set of women who are bank tellers.
- The set of women who are feminists.

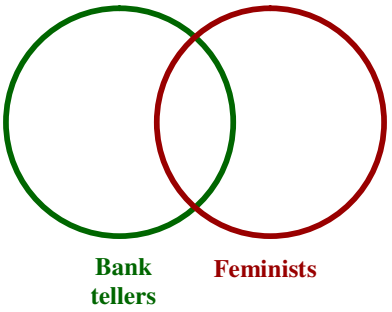

Shade with the pen the area corresponding to women who are bank tellers and feminists.

Appendix B

Stimuli used in Experiment II

Page 1

The Scandinavian peninsula is the European area with the greatest percentage of people with blond hair and blue eyes. This is the case even though (as in Italy) every possible combination of hair and eye color occurs.
Suppose we choose at random 100 individuals from the Scandinavian population. Which group do you think is the most numerous? (Check your choice.)

- Individuals who have blond hair.
- Individuals who have blond hair and blue eyes.
- Individuals who have blond hair and do not have blue eyes.

Page 2

In the picture below there are represented respectively:

- The set of Scandinavian individuals who have blond hair.
- The set of Scandinavian individuals who have blue eyes.

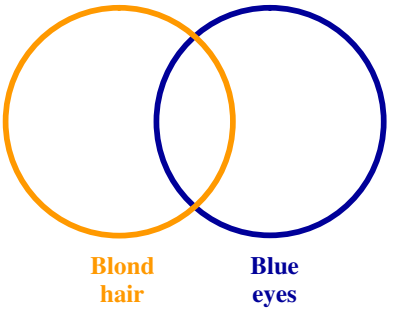

Shade with the pen the area corresponding to Scandinavian individuals who have blond hair and blue eyes.

References

Adam, M. B., & Reyna, V. F. (2005). Coherence and correspondence criteria for rationality: Experts' estimation of risks of sexually transmitted infections. *Journal of Behavioral Decision Making, 18*, 169–186.

Adler, J. (1984). Abstraction is uncooperative. *Journal for the Theory of Social Behaviour, 14*, 165–181.

Agnoli, F. (1991). Development of judgmental heuristics and logical reasoning: Training counteracts the representativeness heuristic. *Cognitive Development, 6*, 195–217.

Ahn, W., & Bailenson, J. (1996). Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *Cognitive Psychology, 31*, 82–123.

Birnbaum, M. H., Anderson, C. J., & Hynan, L. G. (1990). Theories of bias in probability judgment. In J. P. Caverni, J. M. Fabre, & M. Gonzalez (Eds.), *Cognitive biases* (pp. 477–499). North Holland: Elsevier.

Bonini, N., Tentori, K., & Osherson, D. (2004). A different conjunction fallacy. *Mind & Language, 19*, 199–210.

Costello, F. J. (2009). How probability theory explains the conjunction fallacy. *Journal of Behavioral Decision Making, 22*, 171–190.

Crandall, C., & Greenfield, B. (1986). Understanding the conjunction fallacy: A conjunction of effect. *Social Cognition, 4*, 408–419.

Crupi, V., Fitelson, B., & Tentori, K. (2008). Probability, confirmation and the conjunction fallacy. *Thinking & Reasoning, 14*, 182–199.

Davidson, D. (1995). The representativeness heuristic and the conjunction fallacy effect in children's decision making. *Merrill-Palmer Quarterly, 41*, 328–346.

Dulany, D. E., & Hilton, D. J. (1991). Conversational implicature, conscious representation, and the conjunction fallacy. *Social Cognition, 9*, 85–110.

Epstein, S., Donovan, S., & Denes-Raj, V. (1999). The missing link in the paradox of the Linda conjunction problem: Beyond knowing and thinking of the conjunction rule, the intrinsic appeal of heuristic processing. *Personality and Social Psychology Bulletin, 25*, 204–214.

Fisk, J. E. (2004). Conjunction fallacy. In R. F. Pohl (Ed.), *Cognitive illusions: A handbook on fallacies and biases in thinking, judgment, and memory*. London: Psychology Press.

Fisk, J. E., & Slattery, R. (2005). Reasoning about conjunctive probabilistic concepts in childhood. *Canadian Journal of Experimental Psychology, 59*, 168–178.

Frederick, D. M., & Libby, R. (1986). Expertise and auditors' judgments of conjunctive events. *Journal of Accounting Research, 24*, 270–290.

Garb, H. N. (2006). The conjunction effect and clinical judgement. *Journal of Social and Clinical Psychology, 25*, 1048–1056.

Gavanski, I., & Roskos-Ewoldsten, D. R. (1991). Representativeness and conjoint probability. *Journal of Personality and Social Psychology, 61*, 181–194.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A rebuttal to Kahneman and Tversky. *Psychological Review, 103*, 592–596.

Gigerenzer, G. (2001). Content-blind norms, no norms, or good norms? A reply to Vranas. *Cognition, 81*, 93–103.

Gigerenzer, G. (2005). I think, therefore I err. *Social Research, 72*, 1–24.

Hertwig, R., Benz, B., & Krauss, S. (2008). The conjunction fallacy and the many meanings of *and*. *Cognition, 108*, 740–753.

Hertwig, R., & Chase, V. M. (1998). Many reasons or just one: How response mode affects reasoning in the conjunction problem. *Thinking and Reasoning, 4*, 319–352.

Hertwig, R., & Gigerenzer, G. (1999). The 'conjunction fallacy' revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making, 12*, 275–305.

Ho, J. L., & Keller, L. R. (1994). The effect of inference order and experience-related knowledge on diagnostic conjunction probabilities. *Organizational Behavior & Human Decision Processes, 59*, 51–74.

Jarvstad, A., & Hahn, U. (2009). Unreliable sources and the conjunction fallacy. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st annual meeting of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Jones, S. K., Taylor Jones, K., & Frisch, D. (1995). Biases of probability assessment: A comparison of frequency and single-case judgments. *Organizational Behavior and Human Decision Processes, 61*, 109–122.

Kahneman, D., & Frederick, S. (2002). Representativeness revised: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment*. New York: Cambridge University Press.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review, 103*, 582–591.

Koehler, J. J. (1996). The base rate fallacy reconsidered: Normative, descriptive and methodological challenges. *Behavioral and Brain Sciences, 19*, 1–53.

Lagnado, D. A., & Shanks, D. R. (2002). Probability judgment in hierarchical learning: A conflict between predictiveness and coherence. *Cognition, 83*, 81–112.

Lee, M. D., Grothe, E., & Steyvers, M. (2009). Conjunction and disjunction fallacies in prediction markets. In N. Taatgen, H. van Rijn, L. Schomaker, & J. Nerbonne (Eds.), *Proceedings of the 31th annual conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum.

Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.

Mellers, A., Hertwig, R., & Kahneman, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science, 12*, 269–275.

Messer, W., & Griggs, R. (1993). Another look at Linda. *Bulletin of the Psychonomic Society, 31*, 193–196.

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

Moro, R. (2009). On the nature of the conjunction fallacy. *Synthese, 171*, 1–24.

Nilsson, H. (2008). Exploring the conjunction fallacy within a category learning framework. *Journal of Behavioral Decision Making, 21*, 471–490.

Nilsson, H., & Andersson, P. (2010). Making the seemingly impossible appear possible: Effects of conjunction fallacies in evaluations of bets on football games. *Journal of Economic Psychology, 3*, 172–180.

Nilsson, H., Winman, A., Juslin, P., & Hansson, G. (2009). Linda is not a bearded lady: Configural weighting and adding as the cause of extension errors. *Journal of Experimental Psychology: General, 138*, 517–534.

Politzer, G., & Noveck, I. A. (1991). Are conjunction rule violations the result of conversational rule violations? *Journal of Psycholinguistic Research, 20*, 83–103.

Rao, G. (2009). Probability error in diagnosis: The conjunction fallacy among beginning medical students. *Family Medicine, 41*, 262–265.

Shafir, E. B., Smith, E. E., & Osherson, D. N. (1990). Typicality and reasoning fallacies. *Memory & Cognition, 18*, 229–239.

Sides, A., Osherson, D., Bonini, N., & Viale, R. (2002). On the reality of the conjunction fallacy. *Memory & Cognition, 30*, 191–198.

Sloman, S. A., Over, D., Slovak, L., & Stibel, J. M. (2003). Frequency illusions and other fallacies. *Organizational Behavior and Human Decision Processes, 91*, 296–309.

Stolarz-Fantino, S., Fantino, E., Zizzo, D. J., & Wen, J. (2003). The conjunction fallacy: New evidence for robustness. *American Journal of Psychology, 116*, 15–34.

Teigen, K. H., Martinussen, M., & Lund, T. (1996a). Linda versus World Cup: Conjunctive probabilities in three-event fictional and real-life predictions. *Journal of Behavioral Decision Making, 9*, 77–93.

Teigen, K. H., Martinussen, M., & Lund, T. (1996b). Conjunction errors in the prediction of referendum outcomes: Effects of attitude and realism. *Acta Psychologica, 93*, 91–105.

Tentori, K., & Crupi, V. (2009). How the conjunction fallacy is tied to probabilistic confirmation: Some remarks on Schupbach (2009). *Synthese*. doi:10.1007/s11229-009-9701-y.

Tentori, K., Bonini, N., & Osherson, D. (2004). The conjunction fallacy: A misunderstanding about conjunction? *Cognitive Science, 28*, 467–477.

Tentori, K., Crupi, V., & Russo, S. (2011). *On the determinants of the conjunction fallacy: Probability vs. inductive confirmation* (Manuscript).

Tversky, A., & Kahneman, D. (1982). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 84–98). New York: Cambridge University Press.

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

Villejoubert, G., & Mandel, D. R. (2002). The inverse fallacy: An account of deviations from Bayes's theorem and the additivity principle. *Memory & Cognition, 30*, 171–178.

Wedell, D. H., & Moro, R. (2008). Testing boundary conditions for the conjunction fallacy: Effects of response mode, conceptual focus and problem type. *Cognition, 107*, 105–136.

Zhao, J., Shah, A., & Osherson, D. (2009). On the provenance of judgments of conditional probability. *Cognition, 113*, 26–36.

Zizzo, D. J., Stolarz-Fantino, S., Wen, J., & Fantino, E. (2000). A violation of the monotonicity axiom: Experimental evidence on the conjunction fallacy. *Journal of Economic Behavior & Organization, 41*, 263–276.