Logical Inclusion Fallacies – Transfer of Logical Patterns and Noise

Momme von Sydow (momme.von-sydow@uni-heidelberg.de)

Department of Psychology, University of Heidelberg, Hauptstr. 47-51, D-69117 Heidelberg, Germany

Abstract

Traditionally, logical truth has been taken as a necessary criterion for valid predication. As an alternative, a high probability criterion has been proposed. Yet, as long as this criterion makes use of a standard probability interpretation, it does not resolve the problem of inclusion. For instance, the sentence "ravens are black and they can fly" (conjunction) can never be more probable than "Ravens are black or they can fly (or both)" (inclusive disjunction), although it seems irrational that predication of the latter should be more justified than the former. Bayesian Logic (BL) formalizes a pattern-based probability measure for noisy-logical relations that rules out this problem, predicting an entire system of inclusion fallacies (von Sydow, 2009, in press). This proposal aims to uphold the idea of a rational probabilistic basis for predication, while accounting for the fact that extensional probability falls prey to several problems. Here this model is investigated using logical probability judgments after trial-by-trial learning tasks. The results differed from the predictions when modeled on the shown data, but corroborated BL quite well, when based on subjective frequency judgments. Hence, one needs to distinguish the mapping from objective to subjective frequencies and the one from subjective frequencies to subjective (pattern) probabilities. The latter was very well modeled by BL here. Additionally, a transfer and a frequency estimation task provide first evidence for the suggestion that people distinguish not only between logical patterns, but also between observed noise levels.

Keywords: Probability judgments; conjunction fallacy; logic; predication; bias and heuristics; noise levels.

Predication & Narrow Norms of Probability

Traditionally, logical truth and the absence of falsifying cases are taken as nessessary preconditions for a true predication of attributes along with their logical relation. *General* predications, concerned with propositions about classes of entities (e.g., ravens are black AND they can fly), are traditionally only taken to be true if all empirical cases conform to logically allowed (true) cases of a connectives truth-table definition (see Table 1) (Frege, Russell, Whitehead, Wittgenstein, Popper).

The logical truth criterion demands that the proposition hold *deterministically*. That is, in predicate logic, the sentence "*R* are *B* and *F*" would be interpreted as $\forall r\, (B(r) \wedge F(r))$. Taking standard logics as general truth or adequacy criterion of predication, however, involves three fundamental problems:

(1) *The problem of exceptions*: A sentence such as "ravens are black and they can fly" is falsified and refuted by a single evidence. However, even in this example, there

are cases of albino ravens as well as ravens that cannot fly. Formal logics demands "zero tolerance" towards exceptions. Accordingly, "ravens are black and they can fly" would plainly be false. The only correct sentence (corresponding to a logical connective) would be "ravens are black or not, and they can fly or not" (logically referring to the "verum" $B \top F$). For formal purposes, such an understanding may sometimes be helpful, but clearly not as a general norm of predication.

Table 1: Standard truth-table definitions of the conjunction $A \wedge B$, (binary) affirmation A, inclusive disjunction $A \vee B$, and exclusive disjunction $A \gtrless B$.

		$A \wedge B$	A	$A \vee B$	$A \gtrless B$
A	B	true	true	true	false
A	Non-B	false	true	true	true
Non-A	B	false	false	true	true
Non-A	Non-B	false	false	false	false

(2) *The problem of monotonicity*: If predications should be derived inductively, the number of logical subclasses (cells in the truth-table) that need to be assumed to be true can only increase (and never decrease). In this sense, novel evidence can never rule out old evidence. If the first raven is white, even a million black ravens do not allow for the statement "ravens are black." This problem of "inductive monotonicity" is roughly analogous to inferential monotonicity.

(3) *The problem of inclusion*: Even if there are no exceptions to a predication, the extensional interpretation of logic "ravens are black and they can fly" logically entails that "ravens are black *or* they can fly" $(B \wedge F) => (B \vee F)$. Hence, for *any* situation in which one would use an AND-predicate, one would be equally justified to use an OR-predicate. Therefore, if one assumes that predication often aims at providing informative descriptions of whole situations, extensional logic seems not to be suited to such a task. The (extensional) truth-criterion does not penalize overgeneralization.

If one interprets "ravens are black and they can fly" as $P(B \wedge F|R)$, this may help to resolve the problems of exceptions and non-monotonicity (cf. Schurz, 2005). Despite a few albino ravens, the probability remains high. The current proposal aims to uphold the idea of a rational probabilistic basis for predication, while accounting for the fact that standard extensional probability continues to fall prey to the problem of inclusion. The inclusive disjunction, or the cited tautology, has an equally high probability than the plausible conjunction – or in fact a higher one, due to a few albino ravens. The classical conjunction fallacy problem (Tversky & Kahneman, 1983; Kahneman & Frederick, 2002), if based on frequency information, can be seen as a special case of this problem (e.g., Costello, 2005; Lagnardo & Shanks, 2002; cf. von Sydow, 2009, in press). The inclusion rule, in

the conjunction fallacy debate often formulated as the conjunctive rule $P(B \wedge F) \leq P(B)$, has been postulated to be the most simple and fundamental law of probability theory (Tversky & Kahneman, 1983). Likewise, non-standard accounts of probability (e.g., Dempster-Shafer belief functions, Cohen's Baconian probabilities, or fuzzy logic) accept the conjunction rule (cf. Kahneman & Frederick, 2002; Costello, 2005).

In any case, despite several interesting modifying factors (Fiedler, 1988; Hilton, 1995; Gigerenzer, 1996; Hertwig & Gigerenzer, 1999; Mellers, Hertwig, & Kahneman, 2001; Sloman et al., 2003; Bonini et al., 2004; Wedell & Moro, 2008), the phenomenon of the conjunction fallacy proved remarkably stable, even when one used frequency information, ruled out a misunderstanding of hypothesis, and worked with clear set-inclusions (Lagnardo & Shanks, 2002; Tentori et al., 2004; Sides et al., 2002; von Sydow, 2007, 2008; von Sydow, 2007, 2008, 2009, in press). Even the "narrow norm" of the extensional conjunction rule itself has been called into question (note that, despite differences in other regards I follow Gigerenzer here [1996, cf. 2008]).

A rational analysis provides us with a tool to determine the true rational task posed by a situation to an organism (cf. Anderson, 1990). According to the method we must first determine the inferred goal, here the goal of an assertive use of predications whose probability is to be assessed. Predications, in my view, usually aim at providing a description of an overall situation in a single sentence (von Sydow, 2009). Extensional probabilities, in contrast, are concerned with the extension of a specific set. One would need several probability judgments to characterize a full contingency table. Additonally, extensional probability judgments, as we have seen, are quite uninformative for evidence-based predication. $P(B \wedge F) \leq P(B \wedge (F \vee non\text{-}F)) \leq P(B \vee F) \leq P(B \top F)$ or $P(B \wedge non\text{-}F) \leq P(B >< F) \leq P(B \vee F) \leq P(B \top F)$ holds, irrespective of empirical evidence. This yields the absurd result that one always has to prefer the more general statement – at best, the tautology; here, $B \top F$.

Previous accounts to explain frequency-based conjunction fallacies have postulated that people falsely use no probability measure. Representativeness in a first step evokes a prototype (a representative exemplar), and in a second judges the similarity between description and prototypical exemplar (Kahneman & Frederick, 2002). Gigerenzer (1996, 2008) has criticized this concept of being vague. Many other quantitative accounts, however, have been proposed, replacing the target property of probability by another measure, such as inverse probability (cf. Wolford, 1991; Fisk, 1996) or confirmation (increase of probability: Lagnado & Shanks, 2002; Sides et al., 2002; cf. Crupi, Fitelson, & Tentori, 2008).

Bayesian Inductive Pattern Logic

Bayesian Inductive Pattern Logic (*Bayesian Logic*, for short) aims to take the goals of predication seriously and to introduce a pattern-based notion of probability (von Sydow, 2007, 2009). Probability is generally taken to be a polysemous notion (Hertwig & Gigerenzer, 1998; cf. Fiedler, 1988; Wolford, 1991; Fisk, 1996; Sides et al., 2002). Even probability judgments about propositions may demand

different formalizations. Bayesian Logic is a novel, inductive model concerning probabilities of noisy-logical relationships. Here we will assess probabilities, but this specification of probabilities is suitable for a high-probability criterion for predication. These probabilities thus should relate to human probability judgments about logical predications, at least when one is concerned with an overall characterisation of a situation (a pattern). Although a standard Bayesian account alone does not resolve the problem of inclusion (e.g., Fisk, 1996), BL contributes to the renaissance of Bayesian models on the computational level in the cognitive sciences (cf. Oaksford & Chater, 2007; Tenenbaum & Griffiths, 2001, Kruschke, 2008; cf. von Sydow, in press, for demarcation from other models). BL provides posterior probabilities for noisy-logical hypotheses (again specifying patterns of probabilities), given a data pattern D in a 2×2 contingency matrix. Technically, it integrates over different possible noise-levels to achieve a measure for the probability of different logical hypotheses. Although the method may be useful for machine-learning, the psychological question addressed in this paper is: Does this formal model approximate a psychologically important notion of probability in trial-by-trial learning contexts?

BL is based on the Kolmogorov axioms of probability (including the axiom of additivity), applicable only to the inclusion rule on the level of alternative pattern hypotheses rather than on the usual data of subsets. In what follows, the

main steps of the model are sketched, minus the technical details (cf. von Sydow, 2007, 2009, in press):

Figure 1: Contingency table of data, with the three underlying hypotheses regarding the logically described capacity that may have produced the data.

(1) BL assumes that predications that describe situations are normally concerned with propensities, dispositions, or capacities that may have produced the actual observations. Such predications are explanatory constructs underlying the data. Figure 1 represents observed frequencies in a contingency table, as well as three possible hypotheses about underlying probability patterns that can be described in logical terms. For plain relative frequencies, sample size would be irrelevant. According to BL, in contrast, one piece of evidence, for instance, confirming properties A and B, makes the observation "Xs are A and B" less probable than that of a hundred other such cases.

(2) The underlying patterns are treated as alternative hypotheses (even if they overlap or include

Figure 2: Illustration of of PTs with noise-levels r. Dark vs. bright cells represent high vs. low probability.

each other). BL provides probabilities for logical patterns of probabilities as a whole.

(3) A truth-table of a logical connective (cf. Table 1) specifies which cases are allowed or forbidden, but does not specify the probability of allowed cases. Here truth tables can be re-conceived as "logical probability tables" (PTs) via the *assumption of idealization*. We formalzse how ideal logical explanations can account for actual observations. Correspondingly, the *cell* probabilities of what is logically true (in terms of some logical proposition) are assumed to be equally probable (cf. Johnson-Laird, Legrenzi, Girotto, Legrenzi, & Caverni, 1999). Note that the concern is still with deterministic relationships (no randomness or noise [$r = 0$]). The cell probability of a true cell is simply 1 divided by the number of true cells (resulting for instance in $t_{A \wedge B} = 1$, $t_A = 50$, and $t_{A \vee B} = .33$). A false cell has the generative probability of 0 (zero).

(4) Crucially, the underlying logical pattern probabilities need to be formulated in a non-deterministic way – otherwise the tautology would remain most probable in nearly all cases (even when assuming only seldom exceptions). The observed data in Figure 1 cannot be generated by a deterministic *A*-AND-not-*B* pattern. Of the three shown alternative hypotheses, only the "*A or B or both*" hypothesis would apply. As outlined above, we would thus normally be restricted to the use of tautological predicates. Adding the *assumption of uncertainty* to these ideal logical probability tables (cf. Figure 2) allows the formulation of noisy-logical relationships for diverse noise levels. For each logical PT, different degrees of randomness or noise ($0 < r < 1$) are shown (for technical details, see von Sydow, 2007, 2008, 2009, in press). Depending on plausible noise-levels, different explanatory structures may best explain the data.

(5) *Further model steps*. The further model steps involve standard Bayesian statistics and follow from the non-standard assumptions about logical probabilities (cf. von Sydow, in press). First, the likelihood of the obtained data *D*, given any generative connective-noise hypotheses ($P(D|con\text{-}nective\text{-}noise\ hypothesis_i)$) must be specified by a multinomial distribution. Second, Bayes' theorem is used to calculate the posterior probabilities of the hypotheses, given the data: $P(connective\text{-}noise\ hypothesis_i|D)$. In order to calculate the overall probability of a *logical* hypothesis, independent of noise levels *r*, one may sum up the posterior probabilities over all noise levels.

The resulting value is the probability, $P(A\ o_i\ B|D)$, for each of the 16 dyadic logical connectives (o_i is a general placeholder for connectives). The probabilities of different hypotheses are weighted by the posterior probabilities of the noise-levels for each hypothesis (cf. Figure 3). As an illustration, the results for the data vector [7, 3, 2, 0] show that for low noise-levels (or if for noise-levels one uses, rather than a flat prior, one that favors low noise levels) the inclusive disjunction is clearly the most probable (the normative extensional answer). Given preferences for other noise-levels, other hypotheses would become more probable (although extensionally they are always less probable). This nicely illustrates the results of BL (although we will investigate another aspects of noise-levels here).

In contrast to other measures that explain the occurrence

of conjunction fallacies, the resulting measure still concerns a probability – the probability that the situation corresponds to a certain pattern describable in a noisy-logical way, given the observed frequency data. This interpretation of probability, however, obviously differs from standard probabilities. Extensional probabilities, $P_E(A\ o_i\ B|D)$, are concerned with the relative frequency of data in a logically specified subset. Logical pattern probabilities, $P_P(A\ o_i\ B|D)$, as defined by Bayesian Logic, are concerned with the probability of underlying explanatory logical probability patterns, given the data. Thus people who are asked for the probability of the sentence "*X* are *A* and *B*" and who answer $P_P(A \wedge B|X)$ rather than $P_E(A \wedge B|X)$ do nothing wrong, despite the much simpler intentions of the experiment. Psychologically, it is claimed that people who assess overall situations normally use a probability measure roughly approximating noisy-logical pattern probabilities, technically specified by BL.

Figure 3. Illustration of pattern probabilities of logical hypotheses across 11 modelled noise levels *r* for the given data vector [7, 3, 2, 0] (corresponding to the cells of a 2*2 contingence matrix).

Previous tests of this theory have shown that conjunction fallacies do occur in specified situations, even if one uses frequency information, clear set inclusion and rating scales. Moreover, predicted effects of double conjunction fallacies, negations, sample size and pattern sensitivity have been corroborated (cf. von Sydow, 2007, 2008, in press). Furthermore, BL has suggested a generalization of the phenomenon of conjunction fallacies (and disjunction fallacies) to other logical inclusion fallacies using overview formats (von Sydow, 2009).

It will be investigated here whether BL is applicable in trial-by-trial learning situations as well. Alternatively, this particularly "natural" frequency format may lead participants to use extensional probabilities. Here we checked the subjective frequency representations of the participants, and predicted that they were aware of exceptions even when committing inclusion fallacies. Since it is plausible that people do not keep full track of the actual data, but rather reconstruct the input by using the main representational parameters of their representational model, this addresses the question of representation per se. It is suggested that people represent abstract logical propositions about logical combinations of properties and that they have a rough idea of noise-levels. Additionally, it is predicted that people do not

register small differences between cells. Finally, although the sample sizes vary substantially, BL shows that in all cases the sample size is large enough to provide participants with subjectively plausible hypotheses, so that ultimately the sample size is unimportant. Hence it is predicted that participants will strongly underestimate differences in set size. Finally, we use transfer tasks to investigate whether participants transfer the logical connectives (and corresponding inclusion fallacies) and noise-levels within situations.

Experiment

Method

Participants 112 students from the University of Göttingen participated voluntarily in the experiment, receiving either course-credit or 5 Euros as recompense.

Procedure and Materials The computer-based experiment involved a learning phase (Phase 1) and several subsequent tasks: a completion-transfer task (Phase 2), a probability-judgments task (Phase 3), and a frequency-estimation task (Phase 4).

The study employed a scenario resembling, in some aspects, the so-called Linda task (Kahneman & Tversky, 1983). Conjunction fallacies in the Linda-task involved judging P(Linda is a feminist and a bank teller) > P(Linda is a bank teller), after Linda has been described as a feminist. We employed material that allowed use of the same attributes, while replacing "Linda" by "graduates of the Linda school" (i.e., frequency-based instead of single-event probability estimation). At the outset, there was no mention of specific logical hypotheses sought; rather, students were merely informed that they would be answering questions about the schools and the attributes mentioned.

Phase 1 involved six learning-steps, with trial-by-trial information, about graduates of six schools (Table 2). These schools, divided into two groups, were supported by either a workers' foundation (in the example shown: Linda, Burga, and Rike Schools–Type 1) or a Protestant women's organization (Maria, Magda, and Johanna Schools–Type 2). Participants were randomly assigned to four conditions: low vs. high noise, crossed with two other conditions counterbalancing the school-types. For each "condition," the order of schools was also random. For each school, participants were shown data on several graduates, again in random order.

Table 2: Frequencies for conditions of high and low noise Numbers concern the four cells: $f(A \& B)$; $f(\neg A \& B)$; $f(A \& \neg B)$; and $f(\neg A \& \neg B)$.

School	Low noise	High noise	Trials
Linda	18;1;0;0	12;3;2;2	19
Burga	22;1;2;1	12;4;6;4	26
Rike	32;3;3;2	22;7;6;5	40
Maria	0;9;9;1	3;7;7;2	19
Magda	1;12;12;1	3;9;9;5	26
Johanna	2;17;20;1	7;13;15;5	40

Table 2 depicts the frequencies in the high- and low-noise conditions for each school. For each graduate, participants were told whether she was a feminist (or not) and whether she was a bank teller (or not)— by symbols, crossed-out or left intact. For the counterbalancing condition depicted in Table 2, an AND predication is predicted for the worker school-type (the first three listed), and an EITHER-OR predication for the Protestant school-type. Note that even in low-noise schools there are exceptions to the predicted hypotheses; thus, extensionally, other hypotheses should be estimated to be equally if not more probable. (Other theories of CFs do not specify predictions for exclusive vs. inclusive disjunctions).

In Phase 2, two further schools were introduced: Selma (a Type 1 school) and Christina (a Type 2 school). Participants were informed that, due to computer problems, only partial evidence would be available. Furthermore, for each school there were two variants: in the "cell-a" version, participants were given information only on the number of graduates in the a-cell of the contingency table ($f(F \& B) = 11$), with the other fields empty. In the "cell-d" version, the number of graduates were given in the d-cell alone ($f(\neg F \& \neg B) = 3$). Participants were randomly assigned to one of these variants for one school-type and presented to the other variant for the other type. In both cases they were to supply the missing information (completion-transfer task).

In Phase 3 (probability-judgment task), participants determined which of five hypotheses held "most probable" for each school-type: "The graduates of the X-Schools were generally *feminists* (whether bank tellers or not)"; "*bank tellers* (whether feminists or not)"; "*feminists and at the same time bank tellers*"; "*feminists and at the same time no bank tellers*"; "*either bank tellers or feminists*"; or "*feminists who are no bank tellers, or bank tellers who are no feminists, or also feminists and at the same time bank tellers.*" The formulation "and at the same time" was used to rule out CFs based on non-conjunctive meanings of the word "and."

In Phase 4 (frequency-estimation task), participants reproduced "the samples seen" (the observed frequencies) for all six schools, using labelled contingency tables.

Results

Table 3: Percentage and number of hypotheses selected as most probable in the different school-types (collapsed over counterbalancing conditions) and noise conditions.

	C1a AND schools, low noise	C2a AND schools, high noise	C1b XOR schools, low noise	C2b XOR schools, high noise
F	0 % (0)	13 % (7)	4 % (2)	7 % (4)
B	5 % (3)	16 % (9)	5 % (3)	7 % (4)
$F \wedge B$	79% (44)	39 % (22)	5 % (3)	13 % (7)
$F \wedge \neg B$	0 % (0)	2 % (1)	2 % (1)	2 % (1)
$F >< B$	4 % (2)	2 % (1)	66 % (37)	45 % (25)
$F \vee B$	13 % (7)	29 % (16)	18 % (10)	27 % (15)

Table 3 shows the main results for the probability-judgment task for the two school-types (recorded as AND vs. XOR school-types) and the two noise-conditions. Both display inclusion "fallacies"; that is, using extensional probabilities, it would hold that $P_E(F)$, $P_E(B)$ and $P_E(F \vee B)$ are not only equal to but in fact larger than $P_E(F \wedge B)$. Similarly, $P_E(F >< B) < P_E(F \vee B)$. Nevertheless, the results show a dominant proportion of the hypotheses predicted as most

Table 4: Average (and standard error) of the added cell-frequencies and the model-predictions (rounded) for each cell, condition, and the *d*-cell (❶) and *a*-cell (❷) completion tasks. Only the shown frequencies are in italics.

❶	AND, low noise		AND, high noise		XOR, low noise		XOR, high noise	
Cells	Mean (SE)	Pred.	Mean (SE)	Pred.	Mean (SE)	Pred.	Mean (SE)	Pred.
$F \wedge B$	29.1 (4.6)	50	11.5 (2.5)	10	4.3 (1.0)	3	5.3 (0.8)	3
$F \wedge \neg B$	6.5 (0.9)	3	5.8 (1.1)	3	11.6 (1.3)	25	7.8 (0.7)	10
$\neg F \wedge B$	7.7 (1.8)	3	5.9 (0.7)	3	10.8 (1.3)	25	7.8 (0.7)	10
$\neg F \wedge \neg B$	*3; N = 32*	*3*	*3; N = 10*	*3*	*3; N = 18*	*3*	3; N = 11	3

❷	AND, low noise		AND, high noise		XOR, low noise		XOR, high noise	
Cells	Mean (SE)	Pred.	Mean (SE)	Pred.	Mean (SE)	Pred.	Mean (SE)	Pred.
$F \wedge B$	*11; N = 21*	*11*	*11; N = 12*	*11*	*11; N = 19*	*11*	*11; N = 10*	*11*
$F \wedge \neg B$	3.3 (0.6)	1	7.9 (1.5)	5	25.0 (5.3)	50	16.0 (3.0)	30
$\neg F \wedge B$	2.8 (0.4)	1	7.4 (1.1)	5	26.3 (4.8)	50	13.3 (2.7)	30
$\neg F \wedge \neg B$	2.1 (0.4)	1	5.2 (0.6)	5	7.9 (1.4)	11	9.0 (2.3)	11

probable deviating from this norm. The occurrence of these selections can be shown to be clearly above chance-level and to occur in higher proportions than in alternative conditions (C1a and C2a vs. C1b and C2b, cf. Table 3). Yet, the proportion of predicted answers was higher in the low-noise than in the high-noise conditions (C1a vs. C2a, $\chi^2(3, 112) = 17.86$, $p < .001$; C1b vs. C2b, $\chi^2(3, 112) = 5.20$, $p < .05$). Interestingly, the high-noise condition evoked a more frequent selection of inclusive disjunctions. Although this is the extensionally predicted answer, its cause may lie in misrepresentations of subjective frequencies. A tendency to the middle in high-noise conditions may more more easily lead to OR-predications, even if the pattern probabilities are than fully based on BL. (We will return to this later.)

In the completion-transfer task, the overall results did not meet our predications closely, particularly in the high-noise condition (note that space prevents their presentation here). Nevertheless, when focusing on the main group of participants which checked off the predicted hypothesis in the each probability-judgment task respectively (Table 3), Table 4 reflects the predicted results quite closely. That is, many participants transferred both logical patterns and noise levels. These results are ordinally fully consistent with the predictions (based on average modal noise-levels resulting from the *observed* data in the corresponding learning-phases). Additionally, the results suggest that participants were well aware of the existence of exceptions before proceeding to the probability-judgment task, even in the low-noise conditions.

Next, Table 5, presents the mean estimates of subjective frequencies, once again for those subjects who selected the predicted results in the probability-judgment task, and compares this to the frequencies observed. It becomes apparent that the participants' observations reflected both noise levels and logical patterns very well. Moreover, the presence of logical inclusion "fallacies" cannot be dismissed by claiming that participants were unable to indicate extensions or were unaware of exceptions. Participants apparently had little difficulty representing patterns that violated the causal Markov condition and conditional independence of class-attributes which may be represented causally (Rehder, 2003; cf. von Sydow, et al., 2010).

Additionally, the results strongly corroborate the prediction that participants are virtually unaware of the substantial difference in sample-sizes between schools. Participants did also not distinguish between zero observations and low frequency observations, although this would be an essential difference from a falsificationist or logicist perspective.

Table 5: Mean estimated frequencies and observed frequencies for low-noise (❶) and high-noise (❷) conditions, schools, and the four cells (*a*, *b*, *c*, *d*) of a contingency table (collapsed over counterbalancing variables), conditional on the predicted answers in the probability-judgment task

	Estimated frequencies					Observed frequencies				
❶	A	b	c	d	N_{est}	a	b	c	d	N_{obs}
Linda	19	3	3	3	28	18	1	0	0	19
Burga	18	4	4	2	28	22	1	2	1	26
Rike	20	4	4	3	31	32	3	3	2	40
Maria	3	11	11	3	28	0	9	9	1	19
Magda	4	11	11	3	29	1	12	12	1	26
Johan.	3	12	13	3	32	2	17	20	1	40
❷										
Linda	12	6	5	4	26	12	3	2	2	19
Burga	10	5	5	4	24	12	4	6	4	26
Rike	11	5	6	4	27	22	7	6	5	40
Maria	5	7	8	5	26	3	7	7	2	19
Magda	5	8	8	5	26	3	9	9	5	26
Johan.	6	9	7	5	26	7	13	15	5	40

Note: For the first three schools, N = 40 (low noise) and N = 25 (high noise); for the last three, N = 41 (low noise) and N = 22 (high noise)

Table 6: Mean and SE of estimated frequencies for those who selected $F \vee B$ in the probability-judgment task (collapsed over conditions and counterbalancing variables)

	Mean					SE				
❶	a	b	c	d	N_{est}	a	b	c	d	N_{obs}
Linda	8	7	7	4	27	1.0	1.0	1.0	0.6	19
Burga	6	7	8	4	27	1.0	0.9	0.9	0.6	26
Rike	9	9	9	5	34	1.5	1.3	1.3	0.9	40
Maria	8	7	6	3	26	1.0	1.2	0.8	0.4	19
Magda	8	7	6	4	24	1.0	0.9	0.8	0.4	26
Johan.	11	7	7	4	29	1.8	0.7	0.8	0.4	40

Note: For all schools, N = 24.

Finally, Table 6 investigates the largest group deviating from the predictions in the probability-rating task. Here it is assessed whether participants who adopted the *extensionally* correct solution in the probability-rating task (overall 21% selecting inclusive-disjunctions "$F \lor B$") indeed adopted an extensional strategy; also whether their selection could be due to initially misrepresenting frequencies, then applying pattern-probabilities later. Notably, Bayes Logic predicts the selection of inclusive disjunctions as most probable for the averages in all schools—a remarkable statistic, since most participants who selected the extensionally correct solution in fact misrepresented the frequencies, although their strategy was fully coherent with BL.

Discussion

The results corroborate that people use kinds of pattern probabilities and produce inclusion fallacies in line with BL. In the study, these pattern probabilities were exercised despite awareness of large numbers of exceptions. This became clear form frequency estimations. Participants distinguished logical patterns and noise levels in trial-by-trial tasks. Paradoxically, those who apparently adopted a "correct" extensional strategy in fact misrepresented the data—apparently employing pattern probability as well.

Some misrepresentations may have arisen based on the assumption that the two properties *F* and *B* are normally conditionally independent, as suggested by a common-cause interpretation of properties (see Bayes nets formalism: cf. Rehder, 2003). Nevertheless, despite the use of a learning-task, many participants were well able to represent different interaction patterns between attributes on the level of both frequency judgments and (pattern) probability-judgments. Additionally, they could distinguish noise-levels but did not represent sample-size information very well. The ability to learn noisy logical interaction-patterns, including noisy XOR relations, may be explained by the fact that the task concerned merely dyadic property structures (two attributes). Interestingly, people could learn these relations, even with quite noisy data. Ultimately, however, it was only the distinction between subjective and objective frequency representation that permitted the detection of different sources of "error."

One question to be addressed by future research, then, is whether adequate representation of noisy-logical patterns, with corresponding rational probability-judgments in line with Bayes Logic is obtainable in more complex settings as well. Another question is whether the adequate or inadequate representation of subjective frequencies is again mediated by propositional representations or not.

Acknowledgments

References

Costello, F. J. (2005). A unified account of conjunction and disjunction fallacies in people's judgments of likelihood. In *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society* (pp. 494-499). Mahwah, NJ: Laurence Erlbaum Associates.

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

Fiedler, K. (1988). The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research, 50,* 123-129.

Fisk, J. E. (1996). The conjunction effect: Fallacy or Bayesian inference? *Organizational Behavior and Human Decision Processes, 67,* 76-90.

Gigerenzer, G. (2008). *Rationality for Mortals. How people cope with uncertainty.* New York: Oxford University Press.

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

Girotto, V., & Goncales, M. (2001). Solving probabilistic and statistical problems: A matter of information structure and question form. *Cognition, 78,* 247-276.

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

Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin, 118,* 248-271.

Johnson-Laird, P. N., Legrenzi, P., Girotto, V., Legrenzi, S. M., & Caverni, J.-P. (1999). Naive probability: A mental model theory of extensional reasoning. *Psychological Review, 106*, 62-88.

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

Kruschke, J. K. (2008). Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior, 36*, 210-226.

Lagnado, D. A., & Shanks, D. R. (2002). Probability judgment in hierarchical learning: A conflict between predictiveness and coherence. *Cognition, 93*, 81-112.

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

Oaksford, M., & Chater, N. (2007). *Bayesian rationality. The probabilistic approach to human reasoning.* Oxford University Press.

Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29*, 1141–1159.

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

Sloman, St. A., Over, D., Slovak, L., & Stibel, J. M. (2003). Frequency illusions. *Organizational Behavior and Human Processes, 91,* 296-309.

Tenenbaum, J. B. & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*, 629-640.

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

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

Schurz, G. (2005). Non-monotonic reasoning from an evolutionary viewpoint: Ontic, logical and cognitive foundations. *Synthese*, 146, 37-51.

von Sydow, M. (2009). On a General Bayesian Pattern Logic of Frequency-Based Logical Inclusion Fallacies. In N.A. Taatgen & H. van Rijn (Eds.). *Proceedings of the 31st Annual Conference of the Cognitive Science* Society (pp. 248-253). Austin, TX : Cognitive Science Society.

von Sydow, M. (in press). The Bayesian Logic of Frequency-Based Conjunction Fallacies. *Journal of Mathematical Psychology*. doi:10.1016/j.jmp.2010.12.001.

von Sydow, M., Hagmayer, Y., Meder, B. & Waldmann, M. (2010). How Causal Reasoning Can Bias Empirical Evidence. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the Thirty-Second Annual Conference of the Cognitive Science Society* (pp. 2087-2092). Austin, TX: Cognitive Science Society.

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.