

Central Lancashire Online Knowledge (CLoK)

Title	An Evaluation of the Role of Inductive Confirmation in Relation to the Conjunction Fallacy
Type	Article
URL	https://clock.uclan.ac.uk/45674/
DOI	https://doi.org/10.1080/20445911.2023.2182181
Date	2023
Citation	Fisk, John, Marshall, Dean, Rogers, Paul and Stock, Rosemary (2023) An Evaluation of the Role of Inductive Confirmation in Relation to the Conjunction Fallacy. <i>Journal of Cognitive Psychology</i> . ISSN 2044-5911
Creators	Fisk, John, Marshall, Dean, Rogers, Paul and Stock, Rosemary

It is advisable to refer to the publisher's version if you intend to cite from the work.
<https://doi.org/10.1080/20445911.2023.2182181>

For information about Research at UCLan please go to <http://www.uclan.ac.uk/research/>

All outputs in CLoK are protected by Intellectual Property Rights law, including Copyright law. Copyright, IPR and Moral Rights for the works on this site are retained by the individual authors and/or other copyright owners. Terms and conditions for use of this material are defined in the <http://clock.uclan.ac.uk/policies/>



An evaluation of the role of inductive confirmation in relation to the conjunction fallacy

John E. Fisk, Dean A. Marshall, Paul Rogers & Rosemary Stock

To cite this article: John E. Fisk, Dean A. Marshall, Paul Rogers & Rosemary Stock (2023): An evaluation of the role of inductive confirmation in relation to the conjunction fallacy, Journal of Cognitive Psychology, DOI: [10.1080/20445911.2023.2182181](https://doi.org/10.1080/20445911.2023.2182181)

To link to this article: <https://doi.org/10.1080/20445911.2023.2182181>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 28 Feb 2023.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

An evaluation of the role of inductive confirmation in relation to the conjunction fallacy

John E. Fisk^a, Dean A. Marshall^a, Paul Rogers^b and Rosemary Stock^c

^aSchool of Psychology, University of Central Lancashire, Preston, UK; ^bSchool of Humanities and Social Science, University of Brighton, Brighton, UK; ^cSchool of Human and Social Sciences, University of West London, Brentford, UK

ABSTRACT

Inductive confirmation has been proposed as a mechanism giving rise to the conjunction fallacy. For each of five separate vignettes, probability estimates were obtained for a neutral event, for a second event: i.e. the “added conjunct”, and for their conjunction. The added conjunct was selected such that it was inductively confirmed, either by some background evidence provided in the vignette or by the other component event. So as to achieve sufficient statistical power, multilevel models were used to analyse the data. For the added conjunct, the level of confirmation and the posterior probability were significantly associated such that higher levels of confirmation were associated with larger probability estimates. However, there was no significant association between the level of confirmation on the one hand and the incidence of the fallacy and the conjunctive probability on the other.

ARTICLE HISTORY

Received 3 October 2022
Accepted 14 February 2023

KEYWORDS

Conjunction fallacy;
inductive confirmation;
probabilistic judgement;
reasoning

Over the last 50 years or more, a body of evidence has accumulated demonstrating that individuals violate some of the key normative rules of probability theory, for example, Bayes’ theorem (see Pennycook et al., 2022, for a recent review); the disjunction rule (Bar-Hillel & Neter, 1993; Carlson & Yates, 1989; Fisk, 2002); and the conjunction rule (see Fisk, 2022, for a recent review). Probabilistic reasoning plays an essential part in many aspects of our daily routine. The existence of systematic errors is therefore a cause for concern. The present paper focusses on what is considered to be one of the most important probabilistic reasoning errors: the conjunction fallacy. The conjunction fallacy occurs when the probability of the conjunction of two events, $P(A \& B)$, is judged to be more likely than the probability of one or both of the corresponding component events $P(A)$ and $P(B)$, which violates the conjunction rule in probability theory. Nobel prize-winning psychologist and decision theorist Daniel Kahneman has described the extension rule, and the conjunction rule which derives from it,

as “the simplest and most fundamental qualitative law of probability” (Tversky & Kahneman, 1983, p. 294). The fact that it is routinely violated in everyday judgements is therefore of fundamental importance, theoretically and practically.

A simple example demonstrates the conjunction rule. Imagine that a person is selected randomly from the residents of a nursing home. The probability, that person is female *and* over 75 years of age must be no greater than the probability that they are over 75 (since the over 75s in the nursing home will most likely include some males). Similarly, the probability of being over 75 *and* female must be no greater than the probability of being female (since the nursing home will most likely contain some females under 75).

Therefore, judgements such that:

$$P(\text{female} \& \text{over } 75) > P(\text{over } 75) \quad (1)$$

and/or

$$P(\text{female} \& \text{over } 75) > P(\text{female}) \quad (2)$$

violate the conjunction rule and such judgements are considered examples of the conjunction fallacy.

Perhaps the best-known example of the fallacy, devised by Tversky and Kahneman (1983), concerns a hypothetical individual named Linda who is described as follows:

Linda is 31 years old, single, outspoken and very bright, majored in philosophy and, as a student was deeply concerned with issues of discrimination and social justice and participated in anti-nuclear demonstrations.

Having read this description, individuals are then asked to rank various statements in order of their likelihood. Statements can either refer to single propositions (e.g. Linda is a feminist; Linda is a bank teller) or conjunctions of propositions (e.g. Linda is a bank teller *and* is a feminist). When faced with this task, the vast majority of individuals ranked the conjunction (Linda is a bank teller *and* is a feminist) more likely than its component (Linda is a bank teller) which is a conjunction fallacy.

Prominent among current theories of the fallacy is the model proposed by Tentori et al. (2013). Central to their model is the concept of inductive confirmation. Confirmation is defined as the extent to which some new piece of evidence, e , strengthens the credibility of a given hypothesis, h . For example, the evidence might be Linda's description, and the hypothesis, that Linda is a feminist. Although a distinct concept, confirmation is related to probability. That is to say, the degree to which a hypothesis is confirmed by new evidence is a function of the difference between its prior and posterior probabilities, i.e. $P(h)$ and $P(h|e)$. While the precise nature of this functional relationship has not been fully specified, Tentori and co-workers have set out some basic principles. For a given hypothesis, h , and some associated evidence, e , confirmation, as represented by the expression $c(h, e)$, is quantified as follows:

$$\begin{aligned} c(h, e) &= 0 \text{ when } P(h|e) - P(h) = 0 \\ c(h, e) &> 0 \text{ when } P(h|e) - P(h) > 0 \\ c(h, e) &< 0 \text{ when } P(h|e) - P(h) < 0 \end{aligned} \quad (3)$$

Thus the degree of confirmation is dependent on the difference between $P(h|e)$ and $P(h)$. Events are confirmed when this difference exceeds zero and disconfirmed when it is less than zero. Furthermore, increasing levels of confirmation are characterised by ever larger positive values on the measure, greater disconfirmation by more negative values. For example, if the evidence in Linda's

description confirms that she is a feminist then $c(\text{feminist}, \text{evidence}) > 0$ and $P(\text{feminist} | \text{evidence}) > P(\text{Feminist})$.

In applying their concept of confirmation to the conjunction fallacy phenomenon, Tentori et al. first address a prominent feature of the existing literature on the fallacy. They note that typically, one of the component events (e.g. Linda is a bank teller) within the conjunction is judged to be unlikely or seemingly incompatible with the available evidence (Linda's description) while the other component (which they call the added conjunct) is usually consistent with it and therefore judged to be likely (e.g. Linda is a feminist). This combination typically results in the conjunction being assigned a probability greater than that of the unlikely or incompatible event. Thus, while previous accounts have focussed upon the (relatively high) probability of the added conjunct (e.g. Linda is a feminist) in giving rise to the conjunction fallacy, one of the principal findings reported by Tentori et al. (2013) was that the *probability* assigned to the added conjunct did not invariably affect the incidence of the conjunction fallacy. Tentori et al. go on to argue that it is the extent to which the added conjunct (e.g. feminist) is *inductively confirmed* by the background evidence provided in the scenario (e.g. Linda's description) which is the key determinant of whether the fallacy will arise. Broadly speaking, if the added conjunct is confirmed by the evidence provided in the vignette, then the fallacy is more likely to arise. In simple terms, if the added conjunct is perceived as likely because it is confirmed by the evidence (e.g. being a feminist) then it is likely to give rise to the fallacy. If it is perceived to be likely simply because it is associated with a high prior probability (e.g. owning an umbrella) it is less likely to give rise to the fallacy.

In a series of experiments, the Tentori et al. (2013) results were consistent with these expectations. For example, in their violinist scenario, the background evidence described a person as having a degree in violin practice. Participants were more likely to commit the conjunction fallacy when the added conjunct was "gives music lessons" (which is clearly supported by the background evidence) than when the added conjunct was "owns an umbrella". Clearly, as noted above the (relatively high) probability of someone owning an umbrella can be determined by the proportion of people in the general population who own umbrellas. Regarding the alternative added conjunct, "giving music lessons", although this has a low prior probability,

the posterior probability is boosted by the evidence in the vignette (i.e. that the person in question has a degree in violin practice). As predicted, most instances of the conjunction fallacy occurred when the added conjunct was “gives music lessons” rather than when it was “owns an umbrella”.

The key construct in the account proposed by Tentori et al. (2013) is the extent to which the added conjunct is confirmed by the evidence set out in the vignette, e.g. the extent to which the evidence provided, i.e. that “Ollie has a degree in violin performance”, provides confirmatory support to the added conjunct that “Ollie gives music lessons”. Using the Tentori et al. notation, this degree of inductive confirmation is quantified in the following expression $c(h_2, e|h_1)$, where h_2 is the added conjunct (e.g. Ollie gives music lessons), e is the evidence (e.g. Ollie has a degree in violin performance) and h_1 the other conjunct (e.g. “Ollie is an expert mountaineer”). It is worthy of note that in their model, h_1 is usually unrelated to or sometimes disconfirmed by the evidence e .

Tentori et al. also introduce two other measures of inductive confirmation, $c(h_1, e)$ and $c(h_2, h_1|e)$. The first of these represents the extent to which the first conjunct, e.g. “Ollie is an expert mountaineer” is supported or inductively confirmed by the background evidence that “Ollie has a degree in violin practice”. The second, while accepting that the background evidence is a given, represents the extent to which the added conjunct, “Ollie gives music lessons” is supported or inductively confirmed by the other conjunct, e.g. that “Ollie is an expert mountaineer”. In the Ollie example, it would be reasonable to expect that $c(h_2, e|h_1)$ has a positive value while the other two measures, $c(h_1, e)$ and $c(h_2, h_1|e)$, approximate zero or have slightly negative values.

As noted above, Tentori et al. (2013) demonstrate that the incidence of the fallacy is greater when the added conjunct h_2 is inductively confirmed by the evidence e . In general terms, the basis of this effect has been described in Equation (3) above but more specifically they argue that the likelihood of the fallacy is not only an increasing function of $c(h_2, e|h_1)$ but also an increasing function of $c(h_2, h_1|e)$ while a decreasing function of $c(h_1, e)$ (Tentori et al., 2013, p. 248.) Specifically:

$$CF = f[-c(h_1, e), c(h_2, e|h_1), c(h_2, h_1|e)]. \quad (4)$$

Beyond this, Tentori et al. have extended their model to contexts in which there is no explicit

background evidence but where confirmation of the added conjunct is derived by virtue of its relationship with the other component event, specifically $c(h_2, h_1)$. Here, the occurrence of h_1 provides confirmatory support for h_2 and crucially Equation (4) does not apply in this particular context. Thus, Tentori et al. argue that, as an account of the conjunction fallacy, the concept of confirmation applies to both contexts where evidence is provided and where it is not.

To test the Tentori et al. model, two questions are relevant. First, in cases where *evidence, e, is presented*, the question arises as to whether the three measures of inductive confirmation, $c(h_1, e)$, $c(h_2, e|h_1)$, $c(h_2, h_1|e)$, are associated with the incidence of the conjunction fallacy as predicted in Equation (4) above. Second, in cases where *no evidence is provided*, does the incidence of the fallacy increase monotonically with the magnitude of $c(h_2, h_1)$?

Furthermore, it is important to note that the exact process through which the inductively confirmed event results in a higher incidence of the fallacy remains unexplained. The fallacy occurs when the conjunctive event is assigned a higher probability than the less likely component event. Thus, for example, in the context where evidence is provided, either $c(h_2, e|h_1)$ inflates the conjunctive probability $P(h_1 \& h_2|e)$ or somehow reduces the value of $P(h_1)$. *Relative to the normative situation*, it seems implausible that confirmation of the added conjunct (e.g. Ollie gives music lessons) should somehow reduce the probability assigned to the less likely event (e.g. Ollie is an expert mountaineer). Rather, it seems reasonable to assume that the inductively confirmed event somehow inflates the probability of the conjunctive event with higher degrees of confirmation associated with a greater degree of distortion. In any event, one or the other, or both of these two possibilities must be true if the fallacy is to occur. Beyond this, since the incidence of the fallacy is supposedly also a function of the other two confirmation measures, $c(h_1, e)$ and $c(h_2, h_1|e)$, it would be reasonable to expect that these might also exert some influence on the conjunctive probability, $P(h_1 \& h_2|e)$.

Shifting the focus to the other context in which no evidence is provided, the occurrence of h_1 (e.g. getting promoted) provides confirmatory support for h_2 (e.g. getting a pay rise), i.e. $c(h_2, h_1) > 0$. *Relative to the normative situation*, in order for the fallacy to occur, $c(h_2, h_1)$ must somehow inflate the

posterior conjunctive probability, $P(h_1 \& h_2)$, or else somehow reduce one or both component probabilities $P(h_1)$ or $P(h_2)$. The former seems the more plausible explanation.

Lastly, in relation to the added conjunct, what is the relationship between posterior probability and the confirmation measures? For example, how is the probability that Linda is a feminist given the evidence provided in her description, $P(\text{Feminist} | \text{Description})$, related to the degree of inductive confirmation Linda's description confers on the proposition that she is a feminist, $c(\text{Feminist}, \text{Description} | \text{Bank Teller})$. While Tentori et al. do not directly address the exact nature of this relationship, as we note above, they do indicate that in general terms the degree of confirmation, $c(h_2, e | h_1)$, is functionally related to the difference between the prior probability and the posterior probability (given the evidence provided); the larger the latter relative to the former, the greater the degree of inductive confirmation. From that, it seems reasonable to assume that the degree of inductive confirmation, $c(h_2, e | h_1)$, would be significantly related to the posterior probability of the added conjunct $P(h_2 | e)$. Equally the other two confirmation measures, $c(h_1, e)$ and particularly $c(h_2, h_1 | e)$ might also impact the value of $P(h_2 | e)$.

To summarise, the present study seeks to determine:

1. Whether the incidence of the fallacy is functionally related to Equation (4) as predicted by Tentori et al.;
2. How the measures on the right-hand side of Equation (4) affect the magnitudes of $P(h_2 | e)$, $P(h_1 | e)$, and $P(h_2 \& h_1 | e)$;
3. How the incidence of the fallacy and the magnitude of the conjunctive probability are related to $c(h_2, h_1)$ in cases where h_1 provides confirmatory support for h_2 .

Method

Participants

An opportunity sample consisting of eighty participants (34 females) took part in the study. Participants, who were students and staff from a University in the North West of England, ranged in age from 18 to 50 (average 23.05; SD 5.98). As will be argued in the discussion section, a sample of this size has sufficient power to detect a significant effect at $\alpha = .05$.

Materials

Participants estimated probabilities related to five scenarios three of which (Linda, Bill and Tom) were based on those developed by Tversky and Kahneman (1983) and two (Ollie and Rick). based on those from Tentori et al. (2013), single event probabilities, i.e. the added conjunct and the other component event, were as follows: Ollie (music lessons, mountaineer); Linda (feminist, bank teller); Bill (accountant, plays jazz); Tom (over 55, heart attack); and Rick (athletic, under 25). The Ollie, Linda and Bill vignettes included additional evidence in relation to the judgements required. In the Tom and Rick scenarios, no additional evidence was provided.

All participants judged the likelihood of the two single events, h_1 and the added conjunct h_2 , and their conjunction. The order of statements within each vignette was randomised. Participants responded on a 0–100 point scale. Prior to completing the task participants received instructions on the use of the scale. The Ollie vignette is reproduced below.

Ollie has a degree in violin performance.

Now please judge how probable each of the following statements is by entering a number between 0 and 100 for each one:

Statement	Enter a number between 0 and 100
Ollie gives music lessons	
Ollie is an expert mountaineer	
Ollie is an expert mountaineer and gives music lessons	

Participants then completed the inductive confirmation task. For the Ollie, Linda, and Bill scenarios, this involved producing estimates for the key construct $c(h_2, e | h_1)$, and for the other two inductive confirmation measures, $c(h_1, e)$ and $c(h_2, h_1 | e)$. For the other two scenarios, participants produced estimates for $c(h_2, h_1)$. Participants responded using the same visual analogue type scale as that used by Tentori et al. (2013). For each scenario, participants read the relevant information in relation to h_1 , h_2 and where appropriate, e , and judged the degree of inductive confirmation by making a mark on the scale. Instructions and examples were given to participants as to how to use the scale prior to completing this stage of the task. Participants produced separate estimates for each of the scenarios: Ollie, Linda, Bill, Tom and Rick. The Ollie and Tom and Rick scenarios are reproduced below. The labels,

e.g. $c(h_1, e)$, were omitted from the version presented to participants.

Ollie scenario

1. This scenario concerns an individual named Ollie. $c(h_1, e)$

Now consider the following hypothesis (which could be true or false) concerning Ollie:


Ollie is an expert mountaineer.

Now you are given a new piece of information concerning Ollie:

Ollie has a degree in violin performance.

How does the new piece of information that Ollie has a degree in violin performance affect the hypothesis that *Ollie is an expert mountaineer*?

The information that **Ollie has a degree in violin performance**



the hypothesis that *Ollie is an expert mountaineer*.

2. For this next judgement please assume that **Ollie is an expert mountaineer** $c(h_2, e|h_1)$,

Now consider the following hypothesis (which could be true or false) concerning Ollie:


Ollie gives music lessons.

Recall again the piece of information you were given a concerning Ollie:

Ollie has a degree in violin performance.

How does the information that **Ollie has a degree in violin performance** affect the hypothesis that *Ollie gives music lessons*?

The information that Ollie has a degree in violin performance




the hypothesis that *Ollie gives music lessons*.

3. For this next judgement please recall that the information already provided indicates that **Ollie has a degree in violin performance** $c(h_2, h_1|e)$

How does the fact that **Ollie is an expert mountaineer** affect the hypothesis that *Ollie gives music lessons*?

The fact that Ollie is an expert mountaineer



the hypothesis that *Ollie gives music lessons*.

Rick and Tom scenario

Rick and Tom were selected at random from a representative sample of persons responding to a health survey. $c(h_2, h_1)$

Now consider the following two hypotheses (which could be true or false) concerning Rick and Tom:


1. Rick is younger than 25 years old.
2. Tom has had one or more heart attacks.

Now you are given new pieces of information concerning Rick and Tom:

Rick is engaged in athletic competitions. Tom is over 55 years of age

1. How does the new piece of information that Rick is engaged in athletic competitions affect the hypothesis that Rick is younger than 25 years old?


The information that Rick is engaged in athletic competitions



the hypothesis that Rick is younger than 25 years old

2. How does the new piece of information that Tom is over 55 years old affect the hypothesis that Tom has had one or more heart attacks?


The information that Tom is over 55 years old



the hypothesis that Tom has had one or more heart attacks

- How does the new piece of information that that Rick is engaged in athletic competitions affect the hypothesis that Tom has had one or more heart attacks?

The information that Rick is engaged in athletic competitions



the hypothesis that Tom has had one or more heart attacks.

Copies of all vignettes are available from the corresponding author.

Procedure

Participants were briefed as to the nature of the study and verbal consent was obtained. They completed the questionnaire in their own time, in the presence of the experimenter. After the conjunctive probability and inductive confirmation judgements were made, participants produced disjunctive probability estimates for the same set of statements. The interrelationships between the component, conjunctive and disjunctive probability estimates are the subject of a separate study (Fisk et al., 2019). The focus of the present study is the inductive confirmation measures. At the end of the session participants received a written debrief. The study was administered in accordance with the ethical guidelines of the British Psychological Society.

Design

To provide a test of Equation (4) above, the Ollie, Linda and Bill scenarios were amalgamated and generalised linear mixed models analyses were used with the presence or absence of the fallacy as the DV. Following the procedures set out by Tabachnick and Fidell (2007), the binomial distribution was used with logit link function. In addition, linear mixed models analyses were conducted with the conjunctive, more likely and less likely probabilities as separate DVs. In both sets of analyses, the IVs were $c(h_2, e|h_1)$, $c(h_2, h_1|e)$ and $c(h_1, e)$. In order to

model potential violations of sphericity various covariance structures were modelled across the repeated measures component (the scenarios). These included Identity (ID; homogenous variance, zero covariance), Diagonal (DIAG, heterogeneous variance zero covariance), First Order autoregressive (with homogeneous, AR1, and heterogeneous variance ARH1), Compound Symmetry (with homogeneous, CS, and heterogeneous variance, CSH), and Ante dependence (AD1). Unless otherwise noted, the results for the Identity Model (homogeneity of variance and zero covariances) are reported. Where one of the other covariance structures produced a statistically significantly smaller $-2 \text{ Log Likelihood}$ statistic, this will be reported and indicated in the text. For all models reported here, $-2 \text{ Log Likelihood}$ value, the variances and covariances and other relevant parameters are available from the corresponding author.

We also sought to examine the role of confirmation where no background evidence was provided. To this end, the Tom and Rick scenarios were amalgamated and generalised linear mixed models analyses were used with the presence or absence of the fallacy as the dichotomous DV and $c(h_2, h_1)$ as IV. Binomial distribution was used with logit link function. In addition, linear mixed models analyses were conducted with the conjunctive probability as the DV and $c(h_2, h_1)$ as IV.

Additionally, in the mixed models analyses, where possible we modelled the random intercept and slope parameters using the unstructured (UN) model. This makes it possible to establish whether the variance in the intercepts across individual participants relative to the overall average intercept (the constant term in the model equation) differs significantly from zero. Where significant, this indicates that the effect of *factors not captured in the model* differs significantly across participants. Similarly, for each IV, the UN model provides an indication of whether the variance in the slopes across individual participants around the overall average slope parameter (the beta value or parameter estimate for the particular IV in the model equation) differ significantly from zero. Where significant, this indicates that the effect of the IV on the DV differs significantly across participants. These variances populate the diagonal of the UN matrix. Lastly, the UN model produces covariance measures between the intercept and slope for each IV and covariances between the slopes for each pair of IVs. These covariances populate the off-diagonal of

the UN matrix. This allows significant associations between these parameters to be evaluated, for example, whether individual variations in the intercept are significantly associated with individual variations in the slope. Or whether participant-related variations in the slope with respect to one IV are systematically related to variations in the slope for another IV. For example, relative to the overall slopes, do participants who produce shallower slopes with respect to one IV produce steeper slopes with respect to another IV?

Results

Inspection of Table 1 reveals that the incidence of the fallacy ranged between 33% and 44%. The mean inductive confirmation outcomes are also set out in Table 1. As might be expected, the extent to which the added conjunct (e.g. Linda is a feminist) is inductively confirmed by the background evidence (e.g. Linda's description), i.e. $c(h_2, e|h_1)$, is substantially positive for the three relevant scenarios. In the case of the Rick and Tom scenarios, the degree to which the added conjunct is confirmed by the other component, i.e. $c(h_2, h_1)$, is also substantially positive. In all five cases the estimates differed significantly from zero, $p < .001$ in all cases. Thus clearly, as intended, the added conjunct was inductively confirmed. For the scenarios that included the additional evidence, no specific expectation existed for the other two inductive confirmation measures: $c(h_1, e)$ and $c(h_2, h_1|e)$. If anything, it was expected that they might be slightly negative and Table 1 reveals that this was the

case. However, in all but one case, the measures did not differ significantly from zero, $p > .10$ or more, in these cases. In the Bill scenario, the background evidence significantly disconfirmed the other component (*Bill plays jazz*) which at an average of -4.24 was significantly less than zero. Table 1 also contains the posterior probability estimates given the evidence provided. In these cases, the added conjunct, h_2 , was judged to be likely while h_1 was judged to be unlikely. In the Ollie, Linda, and Bill scenarios, the conjunctive probability was, on average, judged to be similar in magnitude although slightly greater than the less likely h_1 component. For the Rick and Tom scenarios, on average, probabilities assigned to all statements were similar in magnitude and judged to be unlikely with the conjunctive event probability exceeding one or both component probabilities.

Mixed models analyses of the fallacy: Ollie, Linda and Bill scenarios

In relation to the three scenarios where additional evidence was presented, the application of generalised linear mixed models allows a direct test of the Tentori et al model (see Equation (4)). The data for the Ollie, Linda and Bill scenarios constituted the level 1 repeated measures component. Participants constituted the level 2 variable. Within a single analysis, this allowed the three inductive confirmation measures $c(h_2, e|h_1)$, $c(h_2, h_1|e)$ and $c(h_1, e)$ to be included as IVs with the presence or absence of the fallacy as the dichotomous DV. In relation to the level 1 component, in all analyses, only the ID

Table 1. Inductive confirmation and probability judgements.

Scenario	Fallacy (%)	Confirmation judgements			Posterior probability judgements		
			Mean	SD		Mean	SD
Ollie	43.75	$c(h_1, e)$	-1.79	11.74	Music Lessons (h_2)	68.36	21.66
		$c(h_2, e h_1)$	17.61	12.26	Mountaineer (h_1)	34.03	21.03
		$c(h_2, h_1 e)$	-2.10	12.06	Conjunction	39.45	27.13
Linda	36.25	$c(h_1, e)$	-1.53	14.33	Feminist (h_2)	70.30	21.80
		$c(h_2, e h_1)$	15.83	15.20	Bank Teller (h_1)	39.51	23.80
		$c(h_2, h_1 e)$	-1.79	12.15	Conjunction	43.69	22.39
Bill	41.25	$c(h_1, e)$	-4.24	17.01	Accountant (h_2)	67.48	21.02
		$c(h_2, e h_1)$	16.11	14.81	Plays Jazz (h_1)	30.46	24.17
		$c(h_2, h_1 e)$	-0.79	13.07	Conjunction	32.78	22.89
Rick	32.50	$c(h_2, h_1)$	13.69	11.60	Under 25 (h_2)	43.65	15.05
					Athletic (h_1)	44.01	17.54
					Conjunction	44.74	19.46
Tom	32.50	$c(h_2, h_1)$	15.59	10.44	Over 55 (h_1)	43.90	15.01
					Heart Attack (h_2)	38.81	18.09
					Conjunction	41.94	17.05

Table 2. Generalised mixed models analyses for the Ollie, Linda and Bill scenarios examining the relationship between the inductive confirmation measures and the presence or absence of the fallacy.

Model/Independent variable	Coefficient	t/Z	p	95% Confidence interval	
				Lower	Upper
Model with no random intercept/ Slope					
Model parameters		t			
Intercept	-0.679	-3.15	.002	-1.103	-0.254
$c(h_1, e)$	-0.025	-2.29	.023	-0.046	-0.003
$c(h_2, e h_1)$	0.013	1.33	.186	-0.006	0.032
$c(h_2, h_1 e)$	0.001	0.09	.927	-0.023	0.025
Model with random intercept/Slopes					
Model parameters		t			
Intercept	-2.237	-2.51	.013	-3.989	-0.484
$c(h_1, e)$	-0.122	-2.59	.010	-0.215	-0.029
$c(h_2, e h_1)$	0.071	1.55	.124	-0.020	0.162
$c(h_2, h_1 e)$	-0.119	-2.02	.045	-0.236	-0.003
Random variances/ covariances:		Z			
Intercept	27.328	3.06	.002	14.391	51.893
Slope $c(h_2, e h_1)$	0.081	3.19	.001	0.044	0.150
Intercept by slope $c(h_2, e h_1)$	-1.094	-2.52	.012	-1.945	-0.243
Slope $c(h_1, e)$	0.068	2.79	.005	0.034	0.138
Slope $c(h_2, h_1 e)$	0.077	2.48	.013	0.035	0.170
Slope $c(h_1, e)$ by $c(h_2, h_1 e)$	-0.008	-0.42	.673	-0.047	0.030

structure converged. The results of the simple model omitting all random factors are reported in the top section of Table 2. It is clear that only one of the inductive confirmation measures, $c(h_1, e)$, was associated with a statistically significant outcome and the negative sign is consistent with the Tentori et al. (2013) model. However, the crucial indicator in the model, $c(h_2, e|h_1)$, did not exert a statistically significant effect, nor did the third component, $c(h_2, h_1|e)$. Overall the model successfully predicted 61.3% of outcomes. Regarding non-fallacious responses, 90.2% were successfully predicted but only 18.6% of fallacy responses were correctly classified.

So as to complete the picture, we repeated the preceding analysis this time including the random intercept and slope terms. The model failed to converge with a single random term so two random terms were included, one with the intercept and the crucial $c(h_2, e|h_1)$ measure, the other with the $c(h_1, e)$ and $c(h_2, h_1|e)$ measures. In relation to the

proportion of fallacies successfully predicted, this dramatically enhanced the performance of the model with 99% of both fallacious and non-fallacious responses now successfully predicted although most of this increase was attributable to the inclusion of the random intercept.¹ The detailed results are shown in the bottom section of Table 2. The crucial measure, $c(h_2, e|h_1)$, remains non-significant while, as in the preceding analyses, $c(h_1, e)$ does register a statistically significant negative effect on the presence or absence of the fallacy. Likewise $c(h_2, h_1|e)$ was just significant but the negative sign is the opposite of what is predicted in the Tentori et al. (2013) model.

Interestingly virtually all random variances and covariances were significant (see the bottom section of Table 2). Importantly, the random intercept term was significant and just as the intercept captures the effects of factors not present in the model (i.e. factors other than the inductive confirmation effects), the random element captures individual differences in the effects of these factors. Regarding the random slope variance for $c(h_2, e|h_1)$, this was statistically significant consistent with the fact that slopes for individual participants differed significantly from the overall non-significant slope. Thus, while for the whole sample, no significant effect was present and although other interpretations are possible, there may be a subset of participants for whom increases in $c(h_2, e|h_1)$ were associated with an increased propensity for the fallacy. It is worthy of note that the covariance between the random slope associated with $c(h_2, e|h_1)$ and the intercept was significant. The negative sign is consistent with the fact that those with less steep slopes, i.e. less sensitive to changes in $c(h_2, e|h_1)$, produced larger intercepts (as noted above the larger intercepts possibly related to factors not captured by the model).²

Mixed models analyses of the conjunctive, more, and less likely posterior probabilities for the Ollie, Linda and Bill scenarios

Given that the fallacy occurs when the conjunctive probability exceeds that of the less likely

¹Inclusion of the intercept on its own resulted in 78.4% of fallacious responses predicted and 90.2% of non-fallacious responses giving an overall success rate of 85.4%.

²The random slope parameter for $c(h_1, e)$ was also statistically significant. According to Tentori et al, this variable is inversely related to the incidence of the fallacy and while it was associated with a significant effect in the main model equation, the statistically significant random slope parameter demonstrates that the strength of this effect varies significantly across participants. In the main model equation, $c(h_2, h_1|e)$ was statistically significant but as noted above the sign was the opposite of that predicted by Tentori et al. However, the fact that the random slope parameter associated with this effect was statistically significant indicates that this aberrant result was not uniform across participants.

component, according to the inductive confirmation account, increases in $c(h_2, e|h_1)$ and $c(h_2, h_1|e)$ and decreases in $c(h_1, e)$ might be expected to inflate the probability assigned to the conjunctive event thereby making the fallacy more likely. So as to examine this expectation, we conducted a linear mixed models analysis with the conjunctive probability as the dependent variable and $c(h_2, e|h_1)$, $c(h_2, h_1|e)$ and $c(h_1, e)$ as independent variables. With regard to modelling the level 1 repeated measures component, the *DIAG* model (with heterogeneous variances and zero covariances) produced the best-fit relative to the *ID* and first-order autoregressive (*AR1*) models. The other models failed to converge. The *DIAG* model failed to converge with a single term representing the random effects so we included two separate terms, the first with the intercept and $c(h_2, e|h_1)$ and the second with $c(h_2, h_1|e)$ and $c(h_1, e)$. As can be seen in the top panel of Table 3, none of the predictors were statistically

significant and as to the random terms only the intercept was statistically significant.

In order to explore the relationship between the posterior probability of the added conjunct and the three confirmation measures, we conducted a linear mixed models analysis with the more likely probability as the dependent variable and $c(h_2, e|h_1)$, $c(h_2, h_1|e)$ and $c(h_1, e)$ as independent variables. The model failed to converge with a single term representing the random effects so we included two separate terms, as was done in the preceding analysis of the conjunctive probability. The results are set out in the middle section of Table 3. As expected, $c(h_2, e|h_1)$ was positive and statistically significant, demonstrating that, for the added conjunct (i.e. the more likely event), higher levels of confirmation were associated with increased probability estimates. Both the random intercept associated with $c(h_2, e|h_1)$ and the random slope variance were statistically significant as was the covariance between them. The negative sign in relation to the covariance indi-

Table 3. The role of the inductive confirmation measures as determinants of the posterior conjunctive, more likely and less likely probabilities in the Ollie, Linda and Bill scenarios.

Dependent variable/Predictors	Coefficient	t/Z	p	95% Confidence interval	
				Lower	Upper
Conjunctive probability					
Model parameters		t			
Intercept	40.884	12.86	<.001	34.504	47.264
$c(h_1, e)$	0.174	1.26	.218	−0.108	0.456
$c(h_2, e h_1)$	−0.166	−1.31	.196	−0.421	0.088
$c(h_2, h_1 e)$	0.025	0.18	.859	−0.267	0.318
Random variances/covariances:		Z			
Intercept	367.25	3.23	.001	200.025	674.293
Slope $c(h_2, e h_1)$	0.266	1.53	.127	0.074	0.961
Intercept by Slope $c(h_2, e h_1)$	−6.953	−1.82	.069	−14.458	0.552
Slope $c(h_1, e)$	0.440	1.65	.100	0.134	1.450
Slope $c(h_2, h_1 e)$	0.142	0.57	.569	0.005	4.426
Slope $c(h_1, e)$ by $c(h_2, h_1 e)$	−0.121	−0.57	.569	−0.536	0.295
More likely probability (the added conjunct)					
Model parameters		t			
Intercept	63.358	21.15	<.001	57.341	69.376
$c(h_1, e)$	−0.051	−0.58	.566	−0.236	0.133
$c(h_2, e h_1)$	0.402	3.52	.001	0.170	0.633
$c(h_2, h_1 e)$	−0.080	−0.57	.614	−0.561	0.401
Random variances/covariances:		Z			
Intercept	432.71	3.83	<.001	259.28	722.16
Slope $c(h_2, e h_1)$	0.388	2.19	.028	0.159	0.949
Intercept by slope $c(h_2, e h_1)$	−10.247	−2.48	.013	−18.361	−2.133
Slope $c(h_1, e)$	0.095	0.87	.384	0.010	0.899
Slope $c(h_2, h_1 e)$	0.583	2.40	.016	0.258	1.321
Slope $c(h_1, e)$ by $c(h_2, h_1 e)$	−0.235	−2.09	.037	−0.455	−0.015
Less likely probability (the other conjunct)					
Model parameters		t			
Intercept	40.089	16.54	<.001	35.297	44.881
$c(h_2, e)$	0.276	3.00	.003	0.095	0.457
$c(h_2, e h_1)$	−0.376	−4.17	<.001	−0.553	−0.198
$c(h_2, h_1 e)$	0.044	0.44	.663	−0.156	0.245
Random variances/covariances:		Z			
Intercept	225.295	4.58	<.001	146.829	345.693

cates that those who perceived a weaker relationship (a flatter slope) between the degree of confirmation and the more likely probability produced larger intercepts relative to those who perceived a stronger relationship. Putting it another way, those who perceived a weaker relationship appear to have based their estimate of the more likely probability on factors other than the evidence provided in the vignette. Neither $c(h_2, h_1|e)$ nor $c(h_1, e)$ were statistically significant as predictors of the more likely component probability.³

For completeness we examined the extent to which the posterior probability of the less likely component, $P(h_1)$, was associated with the confirmation measures. With regard to $c(h_1, e)$, relative to those who perceived no relationship, it might be expected that those who believed that the evidence disconfirmed (confirmed) h_1 might produce smaller (larger) probability estimates for it. We had no expectation in relation to the other two confirmation measures. In order to explore this relationship, we again used linear mixed models analysis. The *DIAG* model performed better than the *ID* with other models failing to converge. None of the models converged when random slope terms were included so only the random intercept term was included. The results are set out in the bottom section of Table 3. Perhaps unsurprisingly, $c(h_1, e)$ was positively associated with the less likely probability $p(h_1|e)$, indicating that those who believed that the evidence did not particularly disconfirm the less likely event or indeed perhaps confirmed it, produced larger estimates of the probability of that event. Rather surprisingly the results revealed that $c(h_2, e|h_1)$ negatively impacted the less likely posterior probability. Specifically, those who believed that the evidence strongly (only weakly) confirmed h_2 , produced smaller (larger) posterior probability estimates for h_1 . As with the preceding analyses of the more likely and conjunctive events, the random intercept term was statistically significant.

Mixed models analyses of the fallacy and of the conjunctive probability for the Tom and Rick scenarios

Shifting attention to the Tom and Rick scenarios, in these cases, no additional evidence was provided

and only one confirmation measure was derived, $c(h_2, h_1)$. We used generalised linear mixed models, combining the data for the two scenarios as the repeated measures level 1 component thereby allowing us to investigate the relationship between the presence or absence of the fallacy and the degree of confirmation, $c(h_2, h_1)$ within a single analysis. Participants constituted the level 2 variable. In relation to the level 1 component, for all analyses, only the *ID* structure converged. The results are set out in the top section of Table 4. The basic model without the random components predicted only 2% of fallacies but 99% of non-fallacies. While the intercept was statistically significant, $c(h_2, h_1)$ was just short of significance. Inclusion of the random intercept and slope terms increased the fit such that 98% of fallacies were now predicted and 99% of non-fallacies. In the full model, all predictors were statistically significant, including $c(h_2, h_1)$. However, as the random slope parameter was also significant, this means that the strength of the positive relationship between the degree of confirmation and the prevalence of the fallacy differed significantly between participants. The random intercept term was significant indicating that the intercepts differed significantly between participants. Also, the slope by intercept random term was significant with the negative sign indicating that those experiencing a stronger confirmation effect on the fallacy produced smaller intercepts.

Continuing with the Tom and Rick scenarios, the bottom section of Table 4 addresses the extent to which higher levels of confirmation were associated with increased levels of the conjunctive probability. We used mixed models analysis with the data for the two scenarios constituting the level 1 component and participants as the level 2. In relation to level 1, the *ID* structure was significantly better than the *Diagonal*, with the *CS* and *CSH* structures failing to converge. Inspection Table 4 reveals that, with the exception of the intercept, none of the predictors nor the random effects were statistically significant. Thus there appears to be no significant relationship between the degree of confirmation, $c(h_2, h_1)$, and the conjunctive probability, $P(h_2 \& h_1)$.

³However, two of the random variances associated with these measures were statistically significant. The random slope for $c(h_2, h_1|e)$ varied significantly from the (non-significant) average slope in the model equation. This raises the possibility that those who judged that the added conjunct was confirmed (or disconfirmed) by the other component produced larger (smaller) estimates of the more likely probability. The significant random covariance between $c(h_2, h_1|e)$ and $c(h_1, e)$ suggests that this possibility was mediated by the extent to which these participants viewed the evidence as confirming or disconfirming the other (less likely) event.

Table 4. The role of the inductive confirmation measure, $c(h_2, h_1)$, in accounting for the fallacy and the conjunctive probability in the Rick and Tom scenarios.

Dependent variable model/Independent variable	coefficient	t/Z	p	95% confidence interval	
				Lower	Upper
Fallacy: present or absent					
Model with no random intercept/slope					
Model parameters		t			
Intercept	−1.151	−3.79	<.001	−1.751	−.552
$c(h_2, h_1)$.028	1.74	.083	−.004	.059
Model with random intercept/slopes					
Model parameters		t			
Intercept	−3.667	−3.655	<.001	−5.649	−1.686
$c(h_2, h_1)$.131	2.19	.030	.013	.250
Random variances/covariances:					
Intercept	33.035	3.387	<.001	18.521	58.925
Slope $c(h_2, h_1)$.116	3.036	.002	.061	.222
Intercept by slope	−1.441	−2.631	.009	−2.514	−.367
Conjunctive probability					
Model with random intercept/slopes					
Model parameters		t			
Intercept	40.378	15.27	<.001	35.049	45.708
$c(h_2, h_1)$.205	1.44	.157	−.083	.492
Random variances/covariances:					
Intercept	215.157	2.67	.008	103.259	448.314
Slope $c(h_2, h_1)$.298	1.04	.300	.045	1.975
Intercept by slope	−2.714	−0.71	.478	−10.215	4.787

Discussion

First, it is important to note that our experimental method produced values for $c(h_2, e|h_1)$ and $c(h_2, h_1)$, which were positive and significantly larger than zero demonstrating that the corresponding concepts were inductively confirmed, as intended. Through generalised mixed models analyses, for the Ollie, Linda, and Bill scenarios we were able to directly test the relationship between the inductive confirmation measures and the presence or absence of the fallacy. With just the inductive confirmation measures, the proportion of successfully predicted fallacies was just 19%. Inclusion of the random intercept term did increase the proportion to 78%, but the intercept term in generalised mixed models analyses is known to reflect the effects of factors not present in the underlying model, i.e. factors other than inductive confirmation. Similarly, the *random* intercept term reflects the effects of individual differences in these factors. Lastly, inclusion of the random slope variances and covariances increased the success rate further to 99%, although, as indicated above, this apparent success rate reflects the effects of factors beyond the scope of the model.

Although $c(h_2, e|h_1)$ was not statistically significant as a predictor, the random slope variance

associated with it was statistically significant. Thus, while the overall slope did not differ significantly from zero, it is possible that for some participants a positive slope was evident. Furthermore, the negative covariance between the random slope and intercept terms suggests that shallower slopes were associated with larger intercepts and vice versa. This is consistent with the possibility that while many fallacies were attributable to other factors, it may be that some fallacious judgements do reflect the effects of the inductive confirmation measures. Indeed, beyond any possible role for $c(h_2, e|h_1)$, there is clear evidence that $c(h_1, e)$ did significantly affect the incidence of the fallacy. In the generalised mixed models analyses, $c(h_1, e)$ reliably exerted a negative effect consistent with the predictions of the Tentori et al. (2013) model. However, the significant random slope term indicates that the strength of this affect was not consistent across participants.

Regarding the Tom and Rick scenarios, inclusion of the random intercept and slope parameters, revealed a statistically significant relationship between the degree of inductive confirmation, $c(h_2, h_1)$, and the prevalence of the fallacy—an outcome consistent with the Tentori et al. model. However, the random slope term was also

significant indicating that the strength of the above-mentioned relationship differed significantly between participants. Moreover, as the random intercept and the negative random slope by intercept interaction were statistically significant, this indicates that, relative to the average, at least some participants are less influenced by the degree of confirmation producing reduced slopes and larger intercepts.

As noted above, in both the Ollie, Linda and Bill analysis and in the Tom and Rick analysis, the random slope term was statistically significant, raising the possibility that at least some participants may have responded in a manner consistent with the Tentori et al. model. To investigate this possibility, it is necessary to examine the relationship between the relative strength of perceived confirmation and the propensity for the fallacy *for each individual*. Specifically, each participant produced five confirmation estimates, three for $c(h_2, e|h_1)$ and two for $c(h_2, h_1)$. While in the majority of cases, the added conjunct was viewed as confirmed, just over half of the participants produced at least some cases where it was unconfirmed/disconfirmed. For these individuals we classified their judgements into three groups, (i) those that were not confirmed (including those disconfirmed), (ii) those weakly confirmed, and (iii) those strongly confirmed.⁴ For each of these three groups, we calculated the mean incidence of the fallacy, for each individual.

The implication of the significant random slope parameter is that the slope characterising the relationship between the degree of confirmation and the incidence of the fallacy differs significantly between participants. Some slopes are clearly consistent with the Tentori et al. model, e.g. $D < W < S$ (where “D” stands for not confirmed/disconfirmed, “W” for weakly confirmed and “S” for strongly confirmed). Furthermore, it is possible to broaden the assumptions underlying the model such that the relationship between the fallacy and the degree of confirmation essentially follows a step-like function with the fallacy being triggered when some threshold level of confirmation is reached. For some, this threshold value

might be associated with weak levels of confirmation, while for others the trigger may be stronger levels of confirmation. This being the case, two other slopes become consistent with the model, $D = W < S$, and $D < W = S$ (subject in all cases to $D = 0$). Other slopes are clearly inconsistent with the model, e.g. $S < W < D$. Examples of model consistent and model inconsistent slopes, produced by participants, are set out in Figures 1 and 2. From the perspective of the null hypothesis, i.e. that no relationship exists between the fallacy and the degree of confirmation, fallacies would be randomly distributed across participants and across the three conditions, D , W , and S . This being the case, given a certain overall prevalence of the fallacy, it is possible to estimate the asymptotic probability of each of the different possible slopes and compare this with the actual outcome. This was done and it was found that the observed frequency (22%) of model consistent slopes was exactly equal to the expected frequency under the null hypothesis, $\chi^2 (df = 1, n = 41) = 0, p = 1.000$.

Many participants viewed all added conjuncts as confirmed and thus did not consider any statements as unconfirmed/disconfirmed. Ignoring the data for disconfirmed statements, it is possible to examine the relationship between the degree of confirmation and the incidence of the fallacy for the whole sample,⁵ focussing solely on weakly and strongly confirmed statements. Four possible slopes exist: $W = S = 0$, corresponding to no fallacies being committed, $W = S > 0$ where the incidence of fallacies is positive and the same under both conditions, $W < S$, consistent with the model and $W > S$, inconsistent. Adopting a broad threshold-based definition of the model, the first two slopes are ambiguous, in the first case it may be that the threshold was not reached, in the second, that a very low threshold obtained in which case they might be viewed as model consistent. As was the case in the preceding paragraph, under the null hypothesis, each slope has a given asymptotic probability of occurring. The actual frequencies of $W < S$ and $W > S$ slopes were 37% and 21%, respectively, compared with an expected frequency of 31% in

⁴Each participant produced five confirmation estimates, three for $c(h_2, e|h_1)$ and two for $c(h_2, h_1)$. For clarity of exposition, here the single notation $c(h_2)$ will be used to refer to both of these. For each participant, these five $c(h_2)$ confirmation estimates were classified into three groups. The first group included statements for which $c(h_2) \leq 0$, i.e., Statements which were *not* confirmed or in fact disconfirmed. For those statements that were confirmed, i.e., Only those for which $c(h_2) > 0$, the mean value for $c(h_2)$ was calculated. Then these statements were split into two groups, those for which $c(h_2) \leq$ the confirmed mean and those for which $c(h_2) >$ the confirmed mean.

⁵In fact 7 participants produced only a single fallacy defined as weakly confirmed and for these seven no slope could be generated.




Figure 1. Model inconsistent slopes for unconfirmed/disconfirmed, weakly confirmed, and strongly confirmed events.




Figure 2. Model consistent slopes for unconfirmed/disconfirmed, weakly confirmed, and strongly confirmed events.

each case.⁶ However, including the other two possible slopes, analysis revealed that, overall, the actual frequencies did not differ significantly from the expected frequencies χ^2 ($df = 3$, $n = 73$) = 4.56, $p = .207$. Examples of model consistent and model inconsistent slopes produced by participants are set out in Figures 3 and 4.

Shifting the focus away from the incidence of the fallacy, surprisingly, and contrary to expectation there was no statistically significant positive relationship between $c(h_2, e|h_1)$ and the posterior conjunctive probability, $P(h_1 \& h_2|e)$. This was evident in the mixed models analysis. Similarly, in relation to the Rick and Tom scenarios, $c(h_2, h_1)$

was not significantly associated with the conjunctive probability. Tentori et al. (2013, pp. 247–248) stated that, in their view, "... The prevalence of the conjunction fallacy is an increasing function of the perceived value of $c(h_2, e|h_1)$ ". A plausible mechanism through which increases in $c(h_2, e|h_1)$ [or in $c(h_2, h_1)$] might increase the likelihood of the fallacy is by directly affecting the magnitude of the conjunctive probability, $P(h_1 \& h_2|e)$ [or $P(h_1 \& h_2)$]. For example, if the conjunctive probability gradually increased monotonically with $c(h_2, e|h_1)$, then eventually it would exceed the less likely probability thereby giving rise to the fallacy and thus increases in $c(h_2, e|h_1)$ would result in an increased incidence

⁶Observed frequencies of slopes $W = S = 0$ and $W = S > 0$ were 29% and 14% compared with expected frequencies of 23% and 15%, respectively.




Figure 3. Model inconsistent slopes for weakly and strongly confirmed events.




Figure 4. Model consistent slopes for weakly and strongly confirmed events.

of the fallacy. However, our results rule this out as a possible explanatory mechanism.

An unexpected outcome emerged in the mixed model analyses of the Ollie, Linda, and Bill scenarios. The results revealed that $c(h_2, e|h_1)$ was statistically significant as a predictor of $P(h_1|e)$. The relationship was a negative one and was in addition to the direct effects of $c(h_1, e)$ which, unsurprisingly, had a positive effect on $P(h_1)$. One explanation for this unexpected relationship might have been that those who viewed the evidence as confirming the added conjunct may have also tended to view it as disconfirming the other component, in which case it would be expected that $c(h_2, e|h_1)$ would be negatively correlated with $c(h_1, e)$. However, while the two measures were negatively correlated in the

Ollie and Bill scenarios, the relationships were short of significance, $p = .184$ and $.052$ respectively. In the Linda scenario, the correlation was positive and non-significant, $p = .253$.

Notwithstanding the results of the present study, it is important to note that the fallacy rate observed by Tentori et al. was greater when the added conjunct was confirmed relative to where it simply had the larger probability. Importantly, Tentori and co-workers maintain that these results are inconsistent with most other accounts of the conjunction fallacy (see for example, Crupi & Tentori, 2016; Tentori & Crupi, 2013; Tentori et al., 2013). Their critique includes some of the more prominent accounts of the fallacy including Nilsson et al.'s (2009) weighted averaging model, Costello and

Watts (2014, 2017) probability theory plus noise model, and Busemeyer et al.'s (2011) quantum probability (QP) model, all of which, they argue, predict that the fallacy should be a function of the *probability* of the added conjunct, with larger probabilities giving rise to a greater likelihood of the fallacy. Looking at these alternative accounts in turn, the key feature of the Nilsson et al. model is suggestion that the probability assigned to the conjunction is the weighted average of the probabilities assigned to the component events, with the less likely event having the larger weight. By construction, given that the weights are assumed to sum to unity, this weighted average will invariably reside in the interval between the component events and thus always give rise to the conjunction fallacy. To account for the fact that some conjunctive judgements are not fallacious, Nilsson et al. proposed that individuals' base estimates were distorted by random noise that sometimes was sufficiently large for the fallacy to be avoided. Since the Nilsson et al. model focusses on the probabilities of the component events, it is hard to see how it can account for the higher rate of the fallacy with confirmed events.

The model proposed by Costello and Watts (2014, 2017) is substantially different. They argue that people reason in a manner consistent with probability theory, but their estimates are subject to a degree of random error. For example, when estimating the probability of an event, A , the individual conducts a search through episodic memory. This search is assumed to involve the application of some classification mechanism which processes each relevant episode to determine whether event A is present, thereby determining the proportion of relevant episodes in which the event occurred. From this, the individual can estimate the probability of event A , i.e. $P_E(A)$. However, this search is error prone and there is a probability, d , that an event will either be missed or believed to have occurred when, in fact, it did not. Joint events, conjunctive or disjunctive, are assumed to be more complex rendering the classification process more difficult resulting in the degree of error being inflated by a small amount, Δd , and so, for example, estimates of conjunctive events $P_E(A \text{ and } B)$, will be subject to a greater degree of error. While these errors may on occasion give rise to biased judgements including the conjunction fallacy, crucially, when the individual estimates are combined into higher level probabilistic concepts,

the errors appear to cancel each other out and the individual's estimates are in line with probability theory. As noted above, Crupi and Tentori (2016) argued that Costello and Watts' model could not account for the fact that conjunctions with a confirmed added conjunct were more likely to give rise to the fallacy than those where the added conjunct simply had the larger probability. For example, Tentori et al. observed that $P(\text{mountaineer and music lessons}) > P(\text{mountaineer and umbrella owner})$, even when $P(\text{umbrella owner}) > P(\text{music lessons})$. However, Costello and Watts (2016a) have shown that their model can accommodate such outcomes under certain circumstances. Furthermore, they have shown that their model is more effective in predicting conjunction fallacy rates relative to that of Tentori et al.

The other prominent model addressed by Tentori and Crupi (2013) is the Busemeyer et al. (2011) QP model. While QP is a complex construct, in the simplest terms the probability of an event (and its negation) is defined by the position of a state vector within a space defined by two orthogonal basis vectors. Where a second event is incompatible with the first its (orthogonal) basis vectors, although sharing a common origin, will be rotated at an acute angle to the basis vectors for the first event. Projections from the state vector to the basis vectors then define the probabilities of the two events. Specifically, the probabilities are equal to the square of the projections. QP is associated with order effects, so for any two events the outcome probabilities are dependent on the order in which the events are evaluated. Crucially, when evaluating the probability of a conjunction (assuming that the two constituent events are incompatible), it has been suggested that the more likely event is evaluated first. This results in a repositioning of the state vector, and the projection from this can be such that the conjunctive probability exceeds the probability of the other less likely event, thereby giving rise to the conjunction fallacy. As they did with Costello and Watts' model, Tentori and Crupi (2013) have argued that the QP model cannot account for the fact that the conjunction fallacy is more prevalent for confirmed added conjuncts relative to other added conjuncts even where the latter are associated with higher probabilities. However, in a comment on the Tentori et al. paper, Busemeyer et al. (2015) demonstrated that, given certain assumptions regarding angle of rotation between the basis vectors and the initial projections from

the first vector sub space, QP can generate results consistent with those derived via inductive confirmation.

Beyond the findings reported above, it is worth noting that confirmation has not been applied to other reasoning fallacies such as conservatism, base rate neglect, and disjunction fallacies. In addition, the adequacy of the Tentori et al. model in providing a comprehensive account of the conjunction fallacy has also been called into question, especially where no explicit evidence is provided. Thus, Jonsson and Assarsson (2016) have argued that the concept of confirmation is not applicable to the reverse conjunction fallacy paradigm in which a property of a whole category is erroneously said to be more likely to be true than it is for a conjunctive subset of that category. For example, individuals believed that the statement: "All sofas have backrests" to be more likely to be true than "All uncomfortable handmade sofas have backrests". While it might be possible to overcome the absence of explicit evidence by modifying the problem format through a process of "evidence mining", Jonsson and Assarsson have demonstrated that attempts to do so either do not explain the reverse fallacy results or else violate the assumptions underpinning the concept of confirmation. Specifically, in the context of the reverse fallacy, they tried applying confirmation with the empty set notionally constituting the evidence, but since the absence of evidence is equivalent to irrelevant evidence both statements receive the same level of confirmation, i.e. no confirmation. They also considered the notion that individuals go beyond the information in the premises searching their prior knowledge so as to construct an evidence base. Other possibilities addressed were that some implicit hypothesis was implied by the statements for which they each provided evidential support or alternatively that the statements might provide evidential support for each other. However, in all cases, Jonsson and Assarsson demonstrated that confirmation was unable to account for the reverse conjunction fallacy.

It has been argued by von Sydow (2016) that a Bayesian approach provides a better account of the conjunction fallacy relative to inductive confirmation. Von Sydow proposes that the conjunction fallacy (and other reasoning fallacies) can be better understood as inclusion fallacies, originating from intensional reasoning processes and Bayesian inference, rather than being attributable to

inductive confirmation. From this perspective, probabilistic constructs can be understood in terms of a two-by-two truth table with the four cells representing the presence (corresponding to a true cell) or the absence (a false cell) of a pair of attributes. Each cell can represent a pair of attribute values, e.g. AB, A not B, B not A, and not A not B. Thus, the top left cell might be defined by a logical connective representing the conjunction of A and B. The truth value of this connective might be represented by a 1 (true) in the top left cell and zeros (false) in the remaining cells. The inclusive disjunction might be represented by three true cells, i.e. in the top two and the bottom left cells. The probability of any true cell within a connective is then defined as one divided by the number of true cells. Tokens (for example hypothetical individuals) may be distributed across these cells. Given any distribution of tokens, it is possible to estimate the probability of each logical connective. For example, if all tokens (without exception) fell within the top left cell. This would imply a probability for the conjunction of 1 which, contrary to extensional probability, would exceed the probability of the inclusive disjunction of 0.33. In essence, unlike extensional probability, the individual is estimating the likelihood of the data (e.g. a given distribution of tokens) given a particular logical connective, e.g. the conjunction of A and B. The probability of the logical connective can then be estimated by applying the Bayes theorem. In this situation, it is perfectly possible for the probability of the conjunctive connective to exceed that of its components and the inclusive disjunction. The situation is rendered more complex by the assumption that first-order probabilities are subject to a degree of random noise characterised by the inclusion of second-order probabilities. Thus, exceptions are permitted with some tokens allowed to reside in false cells. Nonetheless, the application of Bayes theorem can still result in situations where the perceived likelihood of the conjunctive connective is greater than that of the component events and the inclusive disjunction. In two experiments, von Sydow (2016) presented participants with various pattern probability distributions (distributions of tokens) and asked them to select which logical connective was most likely. In comparison with prediction from other accounts of the fallacy, including inductive confirmation, participants selections were more consistent with the Bayesian process. However, von Sydow acknowledges that,

the fact that participants responses were consistent with Bayesian norms does not imply that they were derived by a Bayesian process. Rather, far simpler strategies can account for the results obtained. Furthermore, the distributions of tokens used by von Sydow to test his theory were not representative of the typical conjunction fallacy problem. For example, given the approximate probabilities above for the Bill problem, $P(\text{accountant}) = .67$ and $P(\text{jazz}) = .30$ and assuming conditional independence between $P(\text{accountant})$ and $P(\text{jazz})$, a distribution such that the largest cell was the accountant and not jazz connective would be consistent with the component probabilities with most tokens being associated with the accountant cells. From this, it can be shown that the inclusion fallacy with respect to the conjunctive connective would not occur since $P(\text{jazz}|\text{data}) > P(\text{accountant and jazz}|\text{data})$.

Returning the data reported here and to the evaluation of the Tentori et al. model, some methodological aspects require elucidation. In the Rick and Tom scenarios, the participant made only one confirmation judgement, $c(h_2, h_1)$, and the methodology was the same as that employed by Tentori et al. However, in the Ollie, Linda, and Bill scenarios, it was necessary to augment the methodology used to elicit the inductive confirmation judgements. In the Tentori et al. study, confirmation judgements were limited to the key construct, $c(h_2, e | h_1)$. In order to test their model, it was necessary to amend the methodology so as to obtain estimates of $c(h_1, e)$, $c(h_2, h_1|e)$ as well as $c(h_2, e | h_1)$.

For example, in the version used here, h_1 , "Ollie is an expert mountaineer", is initially posed as a possibility. Next, the evidence, e , is presented: "Ollie has a degree in violin performance". Participants are then asked whether the evidence, e , strengthens or weakens the hypothesis that "Ollie is an expert mountaineer", $c(h_1, e)$. After this, as in the original Tentori et al. version, h_1 is stated as a fact, i.e. "For this next judgement assume that Ollie is an expert mountaineer", then the evidence e , that "Ollie has a degree in violin performance", is *reiterated*, and lastly, as in Tentori et al., the participant is asked whether the evidence e strengthens or weakens the hypothesis that "Ollie gives music lessons". Finally, the roles of h_1 and e are reversed. First, the participant is asked to regard the evidence, e : "Ollie has a degree in violin performance" as a fact, it thereby becoming the conditioning event. The focus is now switched to h_1 "How does the

fact that Ollie is an expert mountaineer affect the hypothesis that Ollie gives music lessons" $c(h_2, h_1|e)$. Thus, with regard to the Ollie, Linda, and Bill scenarios, participants in the present study have to assign different roles to h_1 and e depending on the particular judgement. Examination of the relevant research literature reveals that it is not unusual to ask participants to consider different configurations of this kind, where statements like h_1 and h_2 and the conjunctive combination $h_1 \& h_2$ are simultaneously posed as both possibilities and at the same time as conditioning events, e.g. h_2 given h_1 , and h_1 given h_2 . This has been done sometimes in the context of additional evidence, e , and sometimes not. The results of these studies clearly demonstrate that participants can do this (e.g. Fisher & Wolfe, 2014; Fisk & Pidgeon, 1998; Zhao et al. 2009) as well as consider even more complex possible states as constituted by combinations of different component and joint events and potential conditional relationships involving both affirmation, e.g. $P(A|B)$, and negation, e.g. $P(B|\text{not } A)$ (e.g. Costello & Watts, 2016b). Indeed, when considered holistically, Costello and Watts (2016b) maintain that such complex judgements are essentially normative.

Beyond this, as noted above, in the second judgement the evidence e is *reiterated*. From a purist Bayesian perspective, it might be argued that this makes it "old evidence" and as such it cannot be confirmatory. However, it is now widely acknowledged that old evidence can be confirmatory and numerous modifications to Bayesian epistemology have been proposed to accommodate this (see, for example, Hartmann & Fitelson, 2015; Howson, 1991). The fact that the evidence is reiterated also means that it appears before the participant encounters h_2 whereas in the Tentori et al. version the evidence appears *after* the statements that are to be evaluated. However, it is worthy of note that in most typical examples of conjunction fallacy problems, the evidence e is presented *before* the statements to be evaluated (e.g. Aczel et al., 2016; De Neys et al., 2011; Scherer et al., 2017; Tversky & Kahneman, 1983).

In view of the fact that, for the most part, the results presented here do not lend support for the Tentori et al. model, it must be considered whether or not the sample size was sufficient to detect an effect. Power is difficult to estimate in MLM designs. A number of simulation studies have been conducted to establish the sample size

necessary to provide adequate power in MLM. The key determinant has been found to be the number of level 2 units (i.e. in the present study, participants). It has also been found that in samples exceeding 30 level 2 units, the number of level 1 units (i.e. scenarios in the present study) does not substantially affect power (Ali et al., 2019; Maas & Hox, 2005). Given the sample size in the present study ($N=80$, i.e. 80 level 2 units) the results of the simulation studies set out below suggest that there was sufficient power to detect a significant effect were one present. For example, in Maas and Hox's simulation study, for standard MLM with a continuous DV, estimated regression coefficients for the IVs and the estimated random variances were found to be unbiased (on average, bias was less than 1% for all level 2 group sizes down to 30 level 2 units). Maas and Hox also evaluated the accuracy of the estimated standard errors for the regression coefficients and random variances and the corresponding 95% confidence intervals. Averaged over all simulations, with 50 or more level 2 units, the noncoverage rate (the percentage of occasions when the true parameter estimate was outside the interval) averaged 6% for main effects and interactions and 7% for random variances, levels which Maas and Hox deemed to be acceptable. Using a binary DV, with logit estimation, Paccagnella (2011) found that estimates of regression coefficients for the IVs were unbiased even with as few as 30 level 2 groups. However random variance components were somewhat underestimated even with as many as 320 level 2 groups. However, the degree of underestimation decreased with the number of groups, e.g. with 70 groups the degree of underestimation ranged between 3 and 6%. In addition, with regard to the standard errors and the associated 95% confidence intervals, Paccagnella (2011) found that with 70 level 2 groups, noncoverage rates ranged between 4% and 6% for the regression coefficients but 8 and 10% for the random variance.⁷ Also simulating multilevel logistic regression models, Ali et al. (2019) obtained results broadly consistent with those of Paccagnella (2011). In addition, they found that power (defined as the proportion of times that H_0 of null effect was correctly rejected at a 5 percent level of significance) for main effects and interactions ranged

between 0.78 on average (at 50 level 2 units) and 0.91 (at 100 level 2 units).

To summarise, from the present results, there is little reason to believe that the concept of inductive confirmation plays a significant role in accounting for the conjunction fallacy. The present results suggest that most if not all fallacies are attributable to other factors not addressed in the inductive confirmation account. Even if it had been the case that the fallacy was associated with the degree of confirmation, the lack of a straight forward relationship between $c(h_2, e|h_1)$ and the posterior conjunctive probability leaves unanswered the mechanisms through which such an association is manifested.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data availability statement

The data that support the findings of this study are available from the corresponding author, [JEF], upon reasonable request.

References

- Aczel, B. Szollosi, A., & Bago, B. (2016). Lax monitoring versus logical intuition: The determinants of confidence in conjunction fallacy. *Thinking & Reasoning*, 22(1), 99–117. <https://doi.org/10.1080/13546783.2015.1062801>
- Ali, A., Ali, S., Khan, S. A., Khan, D. M., Abbas, K., & Khalil, A., Manzoor, S., Khalil, U. (2019). Sample size issues in multilevel logistic regression models. *PLoS One* 14(11): e0225427 <https://doi.org/10.1371/journal.pone.0225427>
- Bar-Hillel, M., & Neter, E. (1993). How alike is it versus how likely is it: A disjunction fallacy in probability judgments. *Journal of Personality and Social Psychology*, 65(6), 1119–1131. <https://doi.org/10.1037/0022-3514.65.6.1119>
- Busemeyer, J. R., Pothos, E. M., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, 118(2), 193–218. <https://doi.org/10.1037/a0022542>
- Busemeyer, J. R., Wang, Z., Pothos, E. M., & Trueblood, J. S. (2015). The conjunction fallacy, confirmation, and quantum theory: Comment on Tentori, Crupi, and Russo (2013). *Journal of Experimental Psychology: General*, 144(1), 236–243. <https://doi.org/10.1037/xge0000035>

⁷The higher noncoverage rate for random variances is consistent with underestimation of the standard error which potentially inflates the Type 1 error rate for these estimates. However, it is clear that the random variance estimates set out in Tables 2 and 4 are highly significant. Additionally, random intercept and slope parameters are evaluated on a one tailed basis (Tabachnick & Fidell, 2007) which means that the corresponding p values in the Tables should be further reduced by a half. These lower p values, in and of themselves, reduce the likelihood of a Type 1 error.

- Carlson, B.W. and Yates, J.F. (1989). Disjunction errors in qualitative likelihood judgment. *Organizational Behavior and Human Decision Processes*, 44(3), 368–379. [https://doi.org/10.1016/0749-5978\(89\)90014-9](https://doi.org/10.1016/0749-5978(89)90014-9)
- Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological Review*, 121(3), 463–480. <https://doi.org/10.1037/a0037010>
- Costello F., & Watts, P. (2016a). Probability theory plus noise: Replies to Crupi and Tentori (2016) and to Nilsson, Juslin, and Winman (2016). *Psychological Review*, 123(1), 112–123 <https://doi.org/10.1037/rev0000018>
- Costello, F. & Watts, P. (2016b). People's conditional probability judgments follow probability theory (plus noise). *Cognitive Psychology*, 89(1), 106–133. <https://doi.org/10.1016/j.cogpsych.2016.06.006>
- Costello, F., & Watts, P. (2017). Explaining high conjunction fallacy rates: The probability theory plus noise account. *Journal of Behavioral Decision Making*, 30(2), 304–321. <https://doi.org/10.1002/bdm.1936>
- Crupi, V. & Tentori, K (2016). Noisy probability judgment, the conjunction fallacy, and rationality: Comment on Costello and Watts (2014). (2014). *Psychological Review*, 123(1), 97–102. <https://doi.org/10.1037/a0039539>
- De Neys, W., Cromheeke, S., & Osman, M. (2011). Biased but in doubt: Conflict and decision confidence. *PLoS ONE*, 6(1) e15954. <https://doi.org/10.1371/journal.pone.0015954>
- Fisher, C. R. & Wolfe, C. R. (2014). Are people naïve probability theorists? A further examination of the probability theory + variation model. *Journal of Behavioral Decision Making*, 27(5), 433–443. <https://doi.org/10.1002/bdm.1818>
- Fisk, J. E. (2002). Judgments under uncertainty: Representativeness or potential surprise? *British Journal of Psychology*, 93(4), 431–449. <https://doi.org/10.1348/000712602761381330>
- Fisk, J. E. (2022). Conjunction fallacy. In R. F. Pohl (Ed.), *Cognitive illusions: Intriguing phenomena in thinking, judgment, and memory* (3rd ed.). (pp 27–43) Routledge.
- Fisk, J. E., Marshall, D. A., Rogers, P. & Stock, R. (2019). An account of subjective probability judgment for joint events: Conjunctive and disjunctive. *Scandinavian Journal of Psychology*, 60(5), 405–420. <https://doi.org/10.1111/sjop.12560>
- Fisk, J.E. & Pidgeon, N. (1998). Conditional probabilities, potential surprise, and the conjunction fallacy, *The Quarterly Journal of Experimental Psychology Section A*, 51A(3): 655–681. <https://doi.org/10.1080/713755770>
- Hartmann S., & Fitelson, B. (2015). A new Garber-style solution to the problem of old evidence. *Philosophy of Science*, 82(4), 712–717. <https://doi.org/10.1086/682916>
- Howson, C. (1991). The 'old evidence' problem *The British Journal for the Philosophy of Science*, 42(4), 547–555. <https://doi.org/10.1093/bjps/42.4.547>
- Jonsson, M.L. & Assarsson, E. (2016). A problem for confirmation theoretic accounts of the conjunction fallacy. *Philosophical Studies* 173(2), 437–449. <https://doi.org/10.1007/s11098-015-0500-7>
- Maas, C. J. M. & Hox, J. J. (2005). Sufficient sample sizes for multilevel modeling. *Methodology*, 1(3): 86–92. <https://doi.org/10.1027/1614-2241.1.3.86>
- 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(4), 517–534. <https://doi.org/10.1037/a0017351>
- Paccagnella, O. (2011). Sample size and accuracy of estimates in multilevel models. *Methodology*, 7(3):111–120. <https://doi.org/10.1027/1614-2241/a000029>.
- Pennycook, G., Newton, C., & Thompson, V. A. (2022). Base rate neglect. In R. F. Pohl (Ed.), *Cognitive illusions: Intriguing phenomena in thinking, judgment, and memory* (3rd ed.) (pp 44–60). Routledge.
- Scherer, L. D., Yates, J. F., Baker, S. G., & Valentine, K. D. (2017). The influence of effortful thought and cognitive proficiencies on the conjunction fallacy: Implications for dual-process theories of reasoning and judgment. *Personality and Social Psychology Bulletin*, 43 (6): 874–887. <https://doi.org/10.1177/0146167217700607>
- Tabachnick, B.G. & Fidell, L.S. (2007). *Using multivariate statistics* (5th ed.). Boston: Pearson Education.
- Tentori, K., & Crupi, V. (2013). Why quantum probability does not explain the conjunction fallacy. *Behavioral and Brain Sciences*, 36(3), 308–310. <https://doi.org/10.1017/S0140525X12003123>
- Tentori, K., Crupi, V., & Russo, S. (2013). On the determinants of the conjunction fallacy: Probability versus inductive confirmation. *Journal of Experimental Psychology: General*, 142(1), 235–255. <https://doi.org/10.1037/a0028770>
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315. <https://doi.org/10.1037/0033-295X.90.4.293>
- von Sydow, M. (2016). Towards a pattern-based logic of probability judgements and logical inclusion “fallacies”. *Thinking & Reasoning*, 22(3), 297–335. <https://doi.org/10.1080/13546783.2016.1140678>
- Zhao, J., Shah, A., & Osherson, D. (2009). On the provenance of judgments of conditional probability. *Cognition*, 113(1), 26–36. <https://doi.org/10.1016/j.cognition.2009.07.006>