Normative Accounts of Illusory Correlations

Franziska M. Bott

University of Mannheim

David Kellen

Syracuse University

Karl Christoph Klauer

Albert-Ludwigs-Universität Freiburg

Author Note

Franziska M. Bott and David Kellen contributed equally to this paper. Franziska M. Bott was supported by the DFG grant 2277, Research Training Group "Statistical Modeling in Psychology" (SMiP). Karl Christoph Klauer was supported by DFG Reinhart-Koselleck grant DFG Kl 614/39-1. Scripts and data are made available at https://osf.io/7sdgn/?view_only=5574464a07fd40eaa17d0bfd8a048f1f.

Correspondence: Franziska M. Bott, Department of Psychology, University of Mannheim (f.bott@uni-mannheim.de) or David Kellen, Department of Psychology, Syracuse University (davekellen@gmail.com).

Abstract

When learning about the joint occurrence of different variables, individuals often manifest biases in the associations they infer. In some cases they infer an association when none is present in the observed sample. Other times they infer an association that is contrary to the one that is in fact observed. These *illusory correlations* are often interpreted as being the byproduct of selective processing or as the outcome of an ‘illogical’ pseudocontingency heuristic. More recently, a normative account of illusory correlations has been proposed, according to which they result from an application of Laplace’s Rule of Succession. The present work will discuss the empirical and theoretical limitations associated with this normative account, and argue for its dismissal. As an alternative, we propose a normative account that casts illusory correlations as the expected outcome of a Bayesian reasoner relying on marginal frequencies. We show that this account succeeds in capturing the qualitative patterns found in a corpus of published studies.

Keywords: probability, rationality, biases, illusory correlation, pseudocontingencies

Normative Accounts of Illusory Correlations

Assessments of explanatory sufficiency play an essential role in theoretical debates. At a minimum, any candidate theory or model needs to be able to accommodate an established body of empirical results (i.e., “save the phenomena”, Bogen & Woodward, 1988). Such arguments are particularly impressive when they are associated with normative frameworks, as norms often have limited descriptive power. Take the case of the Kolmogorov axioms of probability (e.g., DeGroot, 1975): people’s choices and judgments often deviate from them, for instance, when the probability of the conjunction of two events (e.g., “*Linda is a feminist and a secretary*”) is judged to be larger than the probability of any single event (e.g., “*Linda is a secretary*”).

Erroneous inferences of this kind are commonly perceived to be “irrational” and have long motivated the development of alternative theoretical accounts that do not conform to probability theory (e.g., *representativeness heuristic*, Tversky & Kahneman, 1974; *configural weighting*, Nilsson, Winman, Juslin, & Hansson, 2009) or modal logic (e.g., *model theory*, Johnson-Laird, Khemlani, & Goodwin, 2015). In fact, the non-negligible failure rates of normative accounts have been one argument of proposals to dismiss normative accounts in favor of more descriptive research programs (see Elqayam & Evans, 2011). Another reaction has been to call for the adoption of extended or alternative normative frameworks (e.g., Buchak, 2013; Busemeyer & Bruza, 2012; Dzhafarov & Kujala, 2016; Oaksford & Chater, 2007; Skovgaard-Olsen, Kellen, Hahn, & Klauer, 2019; Spohn, 2012). Yet another group of researchers has argued that certain normative axioms can be upheld when assuming that they are realized by systems that are subjected to different sources of noise (Costello & Watts, 2014, 2016; Dougherty, Gettys, & Ogden, 1999; Fiedler, 1996; Hilbert, 2012).

In this paper, we will focus on one specific class of effects in probability judgments, which we will broadly refer to as *illusory correlations*. These illusory correlations are shown to occur when individuals attempt to infer the probabilities associated with different

variables based on the observation of skewed samples (e.g., Fiedler, 2000; Hamilton & Gifford, 1976; Kareev, 1995; McConnell, Sherman, & Hamilton, 1994; Sherman et al., 2009). Consider the 2×2 contingency subtable reported at the top of Table 1. This subtable consists of M joint realizations of two binary random variables X and Y , each taking on values 0 and 1. These joint realizations are assumed to be samples $\mathbf{k} = (k_{00}, k_{01}, k_{10}, k_{11})$ taken from a multinomial distribution with the (unknown) probability vector $\mathbf{p} = (p_{00}, p_{01}, p_{10}, p_{11})$. The association between the two variables can be quantified by the Phi coefficient ϕ with $-1 \leq \phi \leq 1$:

$$\phi = \frac{k_{11}k_{00} - k_{10}k_{01}}{\sqrt{k_{1\bullet}k_{0\bullet}k_{\bullet 0}k_{\bullet 1}}} \quad (1)$$

The value $\phi = 0$ indicates that there is no association between X and Y in the sample, whereas positive/negative ϕ values indicate a positive/negative relation between X and Y . Another way of quantifying associations is through the Δp rule (Cheng & Novick, 1990), which compares the probability distribution of a binary random variable conditional on the two possible values of another binary random variable. In the case of variables X and Y in Table 1, we can define

$$\Delta p = P(Y = 0 \mid X = 0) - P(Y = 0 \mid X = 1). \quad (2)$$

As in the case of ϕ , Δp will be zero when there is no association, and positive/negative when there is a positive/negative relationship.

Illusory Correlations

We will classify the illusory correlations discussed into two types, *Type-1* and *Type-S*. This distinction is a direct reference to the idea of Type-1 and Type-S inference errors discussed in the statistical literature (Gelman & Carlin, 2014). Type-1 illusory correlations (IC₁) occur when individuals infer an association between variables when no association

Table 1

General contingency table (with joint frequencies/probabilities k_{ij}/p_{ij} and marginal frequencies $k_{i\bullet}$ and $k_{\bullet j}$) and examples of tables used to elicit Type-1 and Type-S illusory correlations

<hr/> 2 × 2 Contingency Table <hr/>				
		$Y = 0$	$Y = 1$	
	$X = 0$	$k_{00} (p_{00})$	$k_{01} (p_{01})$	$k_{0\bullet}$
	$X = 1$	$k_{10} (p_{10})$	$k_{11} (p_{11})$	$k_{1\bullet}$
		$k_{\bullet 0}$	$k_{\bullet 1}$	
<hr/> Type-1 Illusory Correlation (IC ₁) <hr/>				
		Outcome		
		+	−	
Group	A	27	9	36
	B	9	3	12
		36	12	($M = 48$)
<hr/> Type-S Illusory Correlation (IC _S) <hr/>				
		Outcome		
		+	−	
Group	A	16	8	24
	B	8	0	8
		24	8	($M = 32$)

whatsoever can be found in the observed sample, while *Type-S illusory correlations* (IC_S) correspond to cases in which individuals infer an association that has the *opposite sign* of the one found in the observed sample.

Type-1 Illusory Correlations

First, consider subtable IC₁ in Table 1, which lists the joint and marginal frequencies of variables ‘Group’ and ‘Outcome’. Inspection of the marginal frequencies shows that the majority of the observations comes from Group A, and that ‘+’ is the most common outcome (each majority event is found in 36 of the 48 total instances). Although most observations correspond to outcomes ‘+’ coming from Group A, there is no association between Group and Outcome ($\phi = 0$; note that $P(+|A) = \frac{k_{00}}{k_{0\bullet}} = \frac{27}{36} = .75$ and

$P(+|B) = \frac{k_{10}}{k_{1\bullet}} = \frac{9}{12} = .75$). When people encounter a sample of instances such as the ones summarized in that table, they tend to infer a group-outcome association at the population level, such that the ‘+’ outcome is judged as being more probable under the majority group A than under the minority group B ($\hat{\phi} > 0$ and $\Delta\hat{p} > 0$). As an example, consider Experiment 2 of Hamilton and Gifford (1976): participants were shown a sample of group-outcome instances, with the outcomes being desirable and undesirable behaviors. Group A was the majority group and undesirable behavior the most common outcome (each majority event was found in 24 out of 36 total instances). Although the relative frequency of undesirable behaviors was $\frac{2}{3}$ in both groups, the estimated probabilities of undesirable behaviors in Groups A and B were $\hat{p}_A = .66$ and $\hat{p}_B = .45$, respectively, with $\hat{\phi} = .20$ and $\Delta\hat{p} = .21$.

Pseudocontingency Heuristic

Early proposals of illusory correlations include the differential processing of frequent and infrequent events (e.g., distinctiveness accounts; see Chapman & Chapman, 1969; Hamilton & Gifford, 1976; E. A. Wasserman, Dorner, & Kao, 1990), but have not found considerable empirical support (e.g., Bulli & Primi, 2006; Fiedler, Russer, & Gramm, 1993; Klauer & Meiser, 2000; Meiser, 2003; Meiser & Hewstone, 2001, 2004, 2006). A prominent alternative explanation for the occurrence of illusory correlations is that individuals’ correlational inferences are based on the observed *marginal frequencies* (e.g., Fiedler, Freytag, & Meiser, 2009; Fiedler, Kutzner, & Vogel, 2013). When observing skewed marginal frequencies of two variables, participants are assumed to infer heuristically that the most frequent observations in each variable are associated with each other (e.g., Group A and + in subtable IC₁ of Table 1) as well as the infrequent observations (e.g., Group B and –). This *pseudocontingency heuristic* account is supported by results showing that the inferred associations reflect the skewness found in the marginals, independent of whether joint observations are actually presented (e.g., Fiedler & Freytag,

2004; Meiser, 2006; Meiser & Hewstone, 2004; Meiser, Rummel, & Fleig, 2018; Vadillo, Blanco, Yarritu, & Matute, 2016; Vogel, Kutzner, Fiedler, & Freytag, 2013).

Type-S Illusory Correlations

According to the pseudocontingency account, the inference of an illusory correlation should also not depend on the true contingency between the two focal variables. In line, Fiedler (2010) demonstrated that the inferred associations between two variables were based on the marginal frequencies, independent of their actual contingencies, which were zero (IC_1), consistent with the inference based on marginals, or inconsistent with the inference based on marginals (IC_S). Additionally, consider, for example, the IC_S subtable in Table 1: once again, Group A is the majority group and ‘+’ is the most common outcome. In this case, however, there is a negative group-outcome association ($\phi = -.33$, $\Delta p = -.33$) due to the fact that the proportion of ‘+’ outcomes is larger in Group B ($\frac{8}{8} = 1$) than in Group A ($\frac{16}{24} = .67$). Yet, participants tend to judge the probability of ‘+’ under majority group A to be higher than under minority group B. For example, Meiser et al. (2018, Experiment 1, within-subject condition Casino X) found the average probability estimates to be $\hat{p}_A = .57$ and $\hat{p}_B = .46$, which indicate a positive association ($\Delta\hat{p} = .11$).

Illusory correlations have been shown to increase in strength with decreasing working memory capacity (Eder, Fiedler, & Hamm-Eder, 2011), the more salient marginal frequencies are (Meiser et al., 2018) and the more attention is paid to them (Fleig, Meiser, Ettlin, & Rummel, 2017). Accordingly, many studies on IC_S effects use more complex stimulus distributions including a third (context) variable: Typically, the focal variables’ marginal frequencies are skewed within one context and co-vary across the contexts. By implication, the focal variables are each associated with the context variable. The use of contrasting contexts is assumed to increase the salience of the co-occurrence of skewed marginal frequencies and to thereby increase illusory correlation effects (Fiedler & Freytag,

2004).¹

Two Opposing Accounts of Illusory Correlations

Both types of illusory correlations have been shown to occur across numerous studies using a wide variety of stimulus materials (for reviews, see Fiedler, 2000; Fiedler et al., 2009). Several of these studies involve real-world groups (e.g., occupational groups, Hamilton & Rose, 1980), such that stereotypes are in all likelihood affecting the observed estimates. However, illusory correlations have also been shown to occur with abstract stimuli (e.g., geometric shapes and colors), for which stereotypes are unlikely to exist (e.g., Primi & Agnoli, 2002, Experiment 2). These results suggest that illusory correlations are by and large a byproduct of the general way in which individuals infer associations between sampled variables.

The assumed use of marginal frequencies to infer contingencies is described by some of its proponents as “*logically unwarranted but useful and smart*” (Fiedler et al., 2013, p.328). Simply put, joint frequencies cannot be deduced from marginal frequencies. Yet, the heuristic of associating frequent categories and infrequent categories turns out to have considerable adaptive value: it often succeeds in capturing *actual* contingencies while relying on a simpler data format (Kutzner, Vogel, Freytag, & Fiedler, 2011). Like with many other heuristics, the case made for the inferences based on marginals is couched in the notion that they are highly successful in real-world environments in which they are

¹ As pointed out by a reviewer, it is possible that this association between focal and context variables is behind the observed Type-S illusory correlations. However, there are good reasons to believe that this is not the case. First, we can find the IC_S effect in published studies in which there is no association between focal and context variables. For example, Fiedler (2010) found an IC_S effect in a study in which participants were presented with fictitious students described in terms of four binary variables (**sex**, **city**, **university major**, and **hobby**) and subsequently asked to judge conditional probabilities associated with each pair of variables (e.g., “*What is the % of persons coming from Mannheim, given that they study psychology?*”). Moreover, an unpublished study by the first author found an IC_S effect in an experimental design in which there was no association between focal and contextual variables. In this study, the two contexts (**Machine X** and **Machine Y**) involved two exclusive and distinct pairs of buttons (A/B and C/D, respectively), with one context having a negative button/outcome association ($P(+|A, X) = \frac{18}{24} = .75$ and $P(+|B, X) = \frac{10}{12} = .83$), whereas the other context had no such association ($P(+|C, Y) = \frac{4}{18} = .22$ and $P(+|D, Y) = \frac{4}{18} = .22$). While **Machine Y** trials were only used as filler items, an IC_S effect was found in **Machine X**, with $\Delta\hat{p} = .04$, $t(100) = 2.04$, $p = .044$.

applied as compared to in experiments (i.e., the notion of *ecological rationality*, e.g., Gigerenzer, 2019; Goldstein & Gigerenzer, 2002).

In comparison, a normative account of illusory correlations has recently been proposed by Costello and Watts (2019). Rather than reflecting some heuristic use of marginal frequencies, the occurrence of Type-1 illusory correlations is attributed to the updating of prior information as captured by the *Rule of Succession* proposed by Laplace (1820/1951). When outcome ‘+’ is observed k times in N independent samples coming from a given group, the inferred probability p of ‘+’ occurring is:

$$\hat{p} = \frac{k+1}{N+2}. \quad (3)$$

Applying this rule to the example in subtable IC₁ of Table 1, $\hat{p}_A = \frac{27+1}{36+2} = .74$ and $\hat{p}_B = \frac{9+1}{12+2} = .71$ result as probability estimates, with $\Delta\hat{p} = \hat{p}_A - \hat{p}_B = 0.03$ indicating the presence of an association between group membership and outcomes. Costello and Watts (2019) reported simulations showing that the application of the Rule of Succession to probability estimation outperforms the use of sample proportions or observed relative frequencies in producing Type-1 illusory correlations. Consequently, they argue that Type-1 illusory correlations are in fact the “mathematically correct” response when estimating probabilities based on observed frequencies.

These recent developments present two distinct interpretations of illusory correlations. On the one hand, we have a heuristic account of illusory correlations suggesting that the effects cannot be accommodated within the tenets of probability theory. On the other hand, the newly proposed normative account makes the opposite case, with Type-1 illusory correlations being framed as rational responses that follow the presumably optimal Rule of Succession. In the remainder of this paper, we will argue against these two interpretations. First, we will show that the Rule of Succession has some important limitations, theoretical as well as empirical. Second, we will show that the presumably heuristic pseudocontingency inferences can be seen as a rational consequence of

a Bayesian updating of beliefs when the updates are based on (skewed) marginal frequencies. We will then show that a Bayesian Marginal Model that realizes this theoretical insight is able to accommodate the patterns found in the data at large. Finally, we show that the model is fruitful in generating new predictions, some of which can be tested on existing data and turn out to hold when inspecting previously-published data.

Deconstructing the Rule of Succession

No other formula in the alchemy of logic has exerted more astonishing powers. For it has established the existence of God from total ignorance, and it has measured with numerical precision the probability that the sun will rise tomorrow. (Keynes, 1921, p. 89)

Consider a scenario in which only two events are possible, namely a ‘success’ and a ‘failure’ event. Assume that you encounter k success events in a total of N independent observations. What is the probability that a success will be observed in the $(N + 1)$ th observation? Laplace (1820/1951) showed that, when assuming no prior preference for any probability value, one can reach the solution known as the Rule of Succession:

$$\hat{p} = \frac{\int_0^1 p^{k+1}(1-p)^{N-k} dp}{\int_0^1 p^k(1-p)^{N-k} dp} = \frac{k+1}{N+2}. \quad (4)$$

According to Costello and Watts (2019), the Rule of Succession presents a rationalization of IC₁ effects. Specifically, when asked to judge the probability of a success in a population based on a sample, the Rule of Succession provides a mathematically correct way to estimate the probability p . IC₁ effects are then a byproduct of this rule: the +1 and +2 terms in the numerator and denominator, respectively, will regress the inferred probability towards $\frac{1}{2}$ for each group. The more observations one has of one group, the smaller the deviation from $\frac{k}{N}$, which is caused by the +1 and +2 terms. This means that the estimates for a majority group A will be less regressive towards $\frac{1}{2}$ than estimates for a minority

group B. Hence, the estimates will conform to the inequality $\hat{p}_A > \hat{p}_B > \frac{1}{2}$ when $p > \frac{1}{2}$, or $\hat{p}_A < \hat{p}_B < \frac{1}{2}$ when $p < \frac{1}{2}$. In both cases, the estimates imply a group-outcome contingency when none is present in the data. In other words, Type-1 illusory correlations are predicted. Going beyond, Costello and Watts (2019) developed different computational accounts that implement the Rule of Succession (Algorithms 1 and 2) and showed that they can approximate the IC_1 effects reported by Hamilton and Gifford (1976, Experiment 1) and Van Rooy, Vanhooissen, and Van Overwalle (2013, Experiment 1).

When discussing the Rule of Succession, it is useful to frame it as a specific solution within the broader context of *Bayesian inference* (Jeffreys, 1939): Beliefs are represented by probability distributions that are updated in light of the data using Bayes' Theorem, with the estimated probability $\hat{p}$ corresponding to a *point-summary* of the posterior beliefs. In their presentation of the Rule of Succession, Costello and Watts (2019) state, for instance, that it is not reasonable for participants to “conclude” that $\hat{p} = 0$ when $k = 0$ in a total of $N = 2$ observations (p. 439). Indeed, also when framing the Rule of Succession within the context of Bayesian inference, it would be unreasonable for someone to place all probability mass on 0 given a sample size of $N = 2$. However, under no set of reasonable priors is that the case. The point estimate $\hat{p} = 0$, in whatever way it is computed, is only a *summary* of a distribution that places non-zero probability mass over the unit interval. It does not follow that all other values are deemed impossible or implausible, especially when $N = 2$. Yet, it does not follow, either, that $\hat{p} = 0$ might not be a reasonable point estimate.²

Elaborating on the framing of the Rule of Succession within the context of Bayesian inference, let p be the rate parameter of a Binomial distribution. The parameter p indicates the probability of the occurrence of a ‘success’ event among N independent and

² This statement is a good reflection of Costello and Watts’ (2019) ambiguity regarding the concept of ‘probability’ that they are relying on. The Bayesian foundation of the Rule of Succession explicitly calls for a *subjectivist* interpretation, according to which the strength of our beliefs are represented by probability distributions. However, Costello and Watts’ focus on point estimates and argument against specific values such as $\hat{p} = 0$ speak directly against that. Also of note is the fact that their argument for the rationality of the Rule of Succession (which we will discuss in detail later on) focuses on the *frequentist* properties of $\hat{p}$.

identically-distributed observations. Assuming a prior distribution over p , we can capture a reasoner's preconceptions regarding the probability of a success. We will assume a Beta distribution as a prior given its flexibility in capturing beliefs through its shape parameters α and β . The parameters' values can be interpreted as '*pseudocounts*' in terms of 'successes' and 'failures'. The Beta distribution is a *conjugate prior* of the Binomial distribution. This implies that the posterior beliefs, resulting from the application of Bayes' Theorem when observing k successes out of N observations, will follow a Beta distribution with parameters $\alpha^* = \alpha + k$ and $\beta^* = \beta + N - k$. The mean of this posterior distribution is $\frac{\alpha+k}{\alpha+\beta+N}$. Consequently, when $\alpha = \beta = 1$, we are assuming the prior observations of one success and one failure. This prior distribution is uniform over the $[0, 1]$ interval, indicating no preference for any specific range of probability values. The *posterior mean* under this prior is $\frac{\alpha+k}{\alpha+\beta+N} = \frac{1+k}{2+N}$, which corresponds to the Rule of Succession (see Equation 4).

The shape of the prior distribution will affect the posterior mean. Specifically, the posterior mean is shifted towards the prior mean, an effect known as *shrinkage*. Figure 1 illustrates the updating under different priors such as *Jeffrey's prior* ($\alpha = \beta = \frac{1}{2}$). The amount of shrinkage will depend on the prior distribution as well as on the amount of data. The fewer data the greater the shrinkage will be, or put differently, by virtue of a larger sample size, the data are more successful in overwhelming the prior. The ability of the Rule of Succession to account for Type-1 illusory correlations rests on this fact: Because the minority group B is less often observed than the majority group A, the respective posterior mean for Group B will be closer to the prior mean of $\frac{1}{2}$ (cf. middle panels in Figure 1).

The possibility to vary the prior distribution allows for a generalization of the Rule of Succession. Without loss of generality, let $\alpha = \lambda \cdot \pi_p$ and $\beta = \lambda \cdot (1 - \pi_p)$, with $0 \leq \pi_p \leq 1$ and $\lambda \geq 0$. The posterior mean of p under this parametrization is

$$\hat{p} = \frac{\lambda \cdot \pi_p + k}{\lambda + N}. \quad (5)$$

Parameter λ describes the number of pseudocounts $\alpha + \beta$ implied by the prior and thus, its

Figure 1. Posterior Beta distributions (middle and lower panels) that result from updating the prior distributions (upper panels), parametrized by different shape parameters α and β , when observing the frequencies in Table 1 for Type-1 and Type-S illusory correlations. The vertical lines represent posterior means. While the lower segment lines represent posterior medians, the upper segment lines represent posterior modes.

informativeness. Parameter π_p introduces skew into the prior distribution to the extent to which it deviates from 0.5.³ This generalization of the Rule of Succession corresponds to

³ In the special case of the Rule of Succession, $\alpha = \beta = 1$ and hence, $\lambda = 2$, $\pi_p = .5$.

Carnap's Method of Updating (Carnap, 1952), which has been proposed in the decision-making literature to account for a number of phenomena (e.g., Viscusi, 1989; Wakker, 2002). Viscusi (1989), for instance, showed that its incorporation into the Expected Utility model (Savage, 1954; Von Neumann & Morgenstern, 1944) enables the model to accommodate the paradoxes of Allais. More recently, it has been used to characterize decisions based on experienced options and the ‘description-experience gap’ (Aydogan, 2019; Hertwig, Barron, Weber, & Erev, 2004; see also Wakker, 2002).

The Rational Status of the Rule of Succession

Costello and Watts (2019) motivated the use of the Rule of Succession by arguing that it minimizes errors when estimating the probability of a feature or outcome based on a limited sample. Under the notion that a rational approach is one that minimizes errors, they justify the occurrence of a Type-1 illusory correlation by stating that “*the mathematically correct and hence rational response is to judge the rare feature as more likely in the Minority than the Majority population*” (p. 447). This argument is based on (a) a simulation (Costello & Watts, 2019, Algorithm 1) that computes the posterior mean as an estimate of p given a uniform prior, leading to the conclusion that a “*reasoner who makes this estimate for the underlying probability will be, on average, closest to the true population probability that generated the observed sample*” (p. 440) and on (b) the observation that the use of the posterior mean predicts Type-1 illusory correlations due to shrinkage.

Even when taking for granted that accuracy is a yardstick for rationality, we have to take a closer look at Algorithm 1 they use to establish the rationality of the Rule of Succession, which goes as follows:

1. Set integers k and N , with $k \leq N$.
2. Generate a probability p from a uniform distribution.

3. Generate binomially distributed data with probability p and sample size N .

Whenever k successes occur in the binomial data, record p .

4. Repeat steps 2-3 many times and compute the average of the recorded p s.

The average values coming from this simulation very closely match the values obtained with the Rule of Succession, that is $\frac{k+1}{N+2}$. Costello and Watts (2019) interpret this result as support for the rationality of the rule. However, the simulation algorithm effectively samples from the posterior distribution of p given $\frac{k}{N}$ and a flat prior over p . The average of these posterior samples has to match the Rule of Succession (excusing sampling variability) given that the latter is nothing more than the closed-form expression of the posterior mean of p under the same data and prior. In other words, the two values being compared refer to the exact same quantity.

Moreover, note that the mean of the posterior distribution is only one of many legitimate summary descriptions of a posterior distribution's central tendency, given the observed sample. For instance, it would have been perfectly legitimate to use the *posterior mode* or *maximum a posteriori estimate* (MAP), which in the case of the Beta distribution corresponds to $\frac{\alpha+k-1}{\alpha+\beta+N-2}$. Under the prior parameters $\alpha = \beta = 1$, that yield the Rule of Succession, the posterior mode equals the sample proportion $\frac{k}{N}$, which means that no IC_1 effect would be predicted, an absence that Costello and Watts (2019) consider to be incorrect.

The justification for the use of a given estimator and the discussion of its merits in terms of error minimization requires us to establish how errors are quantified. In other words, it requires a *loss function*. Let $C(\theta, \hat{\theta})$ be a loss function quantifying the discrepancy between a true value θ and its estimate $\hat{\theta}$ (for an overview, see L. Wasserman, 2004, Chap. 12). Among the many possible loss functions, we find the following popular options (with $q > 0$ being a scaling parameter):

- **Quadratic Loss:** $C(\theta, \hat{\theta}) = q \cdot (\theta - \hat{\theta})^2$,

- **Linear Loss:** $C(\theta, \hat{\theta}) = q \cdot |\theta - \hat{\theta}|$,
- **All-or-Nothing Loss:** $C(\theta, \hat{\theta}) = \begin{cases} 0, & \text{if } \theta = \hat{\theta}, \\ q, & \text{otherwise.} \end{cases}$

Under a quadratic loss function, the optimal estimator is the posterior mean. The only case in which illusory correlations are not expected using posterior means is when the completely uninformative, so-called *Haldane prior* ($\alpha = \beta = 0$) is used. In the case of linear loss, the optimal estimate corresponds to the posterior median, which is approximated by $\frac{\alpha+k-1/3}{\alpha+\beta+N-2/3}$ (Kerman, 2011). If $\alpha = \beta < \frac{1}{3}$, a *reversed* IC_1 effect is predicted, such that the estimated posterior median of the majority group is closer to $\frac{1}{2}$ than the minority group's (.751 and .754 respectively, in the IC_1 example in Table 1 when $\alpha = \beta = \frac{1}{4}$). Finally, an all-or-nothing loss function is optimized by the posterior mode, which predicts a reversed IC_1 for $\alpha = \beta < 1$ and no IC_1 effect under the uniform prior $\alpha = \beta = 1$. Type-1 illusory correlations are only expected when $\alpha = \beta > 1$. Based on these cases, we see that the Rule of Succession and the illusory correlations that it can predict are only rational/optimal if one takes for granted that $\hat{p}$ corresponds to the posterior mean and a quadratic loss function is assumed. In fact, the estimator $\hat{p}$ deemed optimal differs according to the assumed prior and loss function, with illusory correlations being expected under some but not others.

The optimality results above describe a reasoner who moves in an ecology in which s/he randomly encounters N samples from groups with different true probabilities p that follow the assumed prior distribution. In this context, optimality refers to a minimization of the average estimation error *across* groups. This, however, is at odds with the conceptual characterizations that Costello and Watts (2019) allude to. They conceptualize the prior and posterior distributions as epistemic probability distributions or as descriptions of the reasoner's prior and posterior beliefs, respectively, about the unknown value p underlying the observed group frequencies. Accordingly, the group in question is *fixed* and the true probability that generates the data is the unknown p value in force for

that group. Let us call this fixed value p_G . The uniform prior describes the assumption that in the absence of data, the reasoner considers each admissible value for the unknown true p_G value as equally likely a priori. An estimate $\hat{p}_G$ is in error to the extent to which it deviates from the true value p_G of the group under scrutiny, irrespective of one's prior beliefs. The expected estimation error or loss must be computed as what is to be expected when sampling N observations from group G , that is relative to the binomial distribution with parameters N and p_G . If the sample proportion $\frac{k}{N}$ is taken as estimate, then the expected quadratic loss corresponds to its variance under the binomial distribution, namely $\frac{1}{N}p_G(1 - p_G)$. In the case of the Rule of Succession, $\frac{k+1}{N+2}$, the expected quadratic loss is $\frac{N^2}{(N+2)^2} \frac{1}{N}p_G(1 - p_G) + \frac{1}{(N+2)^2}(2p_G - 1)^2$. The latter expectation is larger (i.e, worse) than the former when $|p_G - \frac{1}{2}| > \sqrt{\frac{N+1}{8N+4}}$, and smaller otherwise. What this result means is that the question of which of these estimators is the best depends on both p_G and N and cannot be answered in the absence of knowledge about p_G .

The Empirical Status of the Rule of Succession

Even if one cannot make a compelling case for the Rule of Succession or its generalization à la Carnap, theoretically, they might nevertheless succeed in capturing the effects found in the data at large. In this section, we will show that this is not the case.

The Case of Type-S Illusory Correlations. Costello and Watts (2019) discussed the Rule of Succession within the context of IC_1 effects. However, there is no reason to not consider IC_S effects as well, especially given that the difference between both types can be boiled down to the relative frequencies observed in each group (namely, whether they are the same or not). It turns out that the Rule of Succession has trouble capturing IC_S effects: First, consider the simulation reported in Figure 2 that displays the probability inferred by the Rule of Succession as a function of the number of observations across the different groups. Note that in order to illustrate these predictions, we replaced k with $N \cdot p$ (see also Viscusi, 1989; Wakker, 2002). The top panel of Figure 2 shows that the

Figure 2. Probability of success predicted by Rule of Succession as a function of number of observations (N) per option.

Rule of Succession generally expects IC₁ effects to occur. But as shown in the center and bottom panels, an IC_S effect is only expected under the Rule of Succession for very specific frequencies, and/or extremely small samples (e.g., when $P(+|A) = .75$, $P(+|B) = .83$, and $N_A < 4$). Otherwise, according to the Rule of Succession, the estimate $\hat{p}$ is expected to be higher for the minority group(s), in line with the genuine group-outcome relationship.

Qualitative Misses in Illusory Correlations. The ability to capture an *effect*, a difference between two estimated values, does not imply the ability to accommodate the values from which the effect is computed. Considering the absolute values of people’s probability judgments more closely, it turns out that beyond the result $\hat{p}_A > \hat{p}_B$, data on IC_1 and IC_S effects show a recurring pattern of $\hat{p}_B < \frac{1}{2}$: In other words, the majority event for Group B is in fact perceived as the minority event for that group. These recurring qualitative misses cannot be explained by any regression-to-the-mean account with prior mean $\frac{1}{2}$, as they are unable to cross over the $\frac{1}{2}$ -midpoint. We checked the occurrence of these qualitative misses in a corpus of previously-published studies, which are listed in Table 2.

The study corpus was built by searching the databases *PSYNDEX* and *PsycInfo* (on 08/20/2019) using the terms “illusory correlation” and “pseudocontingencies” and restricting the search to manuscripts published in academic journals in the period 2000-2019. Of the articles identified, we only considered experiments that additionally met the following criteria: (1) the experimental stimuli consisted of binary variables, (2) participants had to judge the frequency or probability of an outcome given one group, (3) the groups were artificial, without any stereotypical information associated with them, (4) the frequency estimates or probability estimates are reported and were not converted into a Δp measure, (5) the effect addressed (i.e., IC_1 or IC_S) was present in the data and determined to be reliable using some statistical test,⁴ and (6) only non-clinical participants took part in the study. The application of those inclusion criteria amounted to a set of 43 experimental conditions in 14 articles. We additionally included data from Fiedler et al. (1993), Spears, van der Pligt, and Eiser (1985), and Hamilton and Gifford (1976).⁵ The aim was not to provide an exhaustive literature review, but rather to illustrate the general

⁴ In several cases, application of this criterion resulted in including only one of two within-subject conditions.

⁵ Note that applying the stated criteria resulted in not including the two data sets used by Costello and Watts (2019).

Figure 3. Observed and estimated relative frequencies in the majority (A) and minority (B) groups in the data corpus. In order to improve the discriminability between data points, we jittered the observed proportions.

pattern of empirical findings in contemporary research.

Checking the data corpus for the occurrence of qualitative misses described above ($\hat{p}_B < \frac{1}{2}$), we found these misses in 55% of the cases. The magnitude of the qualitative misses is additionally displayed in the right panel of Figure 3: in more than half of the cases, participants' estimates for the minority group do not only reflect an underestimation, but a reversal in which outcome is most frequent. Importantly, note that the qualitative misses are found in both IC_1 and IC_S studies, as can be seen in Table 2, so that they cannot be attributed to the different designs used to study Type-1 and Type-S illusory correlations (see Footnote 1).

Fitting the Rule of Succession to the Data Corpus. To corroborate the above assessments, we fitted the Rule of Succession as well as Carnap's generalization to the corpus of studies reported in Table 2. In the case of the Rule of Succession, for each study, we applied the Rule of Succession to the observed frequencies k and N for both groups A and B. In the case of Carnap's generalization, we fixed π to $\frac{1}{2}$ and estimated λ freely. For each study, we assumed a single λ parameter for both Group A and Group B when

computing estimates. Models were estimated by minimizing squared errors, with the constraint that IC_S and qualitative misses were to be predicted if they were present, and the model able to produce them. These fits tell us how well the models can capture the qualitative patterns in the data, irrespective of their magnitude.⁶ The best-fitting predictions are illustrated in Figure 4.

It is clear that the Rule of Succession cannot provide a reasonable account of the data, especially when it comes to IC_S effects and qualitative misses.⁷ The fact that it has no free parameters also compromises the model’s ability to successfully capture differences between studies in which participants observed the same joint events. The generalized rule performs somewhat better, although some of the effects are barely captured at the expense of extreme prior parameters (e.g., $\alpha = \beta = 50$).⁸

Interim Discussion

The case for the Rule of Succession as a successful normative account stands on two legs. The first one is the notion that it constitutes an optimal solution to an inference problem. The second is its ability to capture the observed effects in the data. What our

⁶ A purely quantitative comparison between observed and predicted values is not really useful here as we first and foremost want to know whether the model can capture the qualitative patterns (reproducing a qualitative pattern is not guaranteed when simply minimizing misfit). Also, one should keep in mind that these data are an amalgamation of different strategies and beliefs that are unlikely to be successfully captured by any single model. In this situation it is best to focus on the qualitative aspects of models (for discussions, see Kellen, 2019; Navarro, 2019).

⁷ As a sanity check, we investigated the possibility that qualitative misses may be caused by sampling noise and/or forgetting. We simulated probability estimates produced by the Rule of Succession including forgetting rates for the two example distributions in Table 1 and the less extreme IC_S -event distribution from Meiser et al. (2018, Experiment 3). Following Algorithm 2 of Costello and Watts (2019), each observation is recalled with probability $1 - f$. This is achieved by comparing a number randomly drawn from a uniform distribution between 0 and 1, against a fixed value f . If the random number is larger than f , the observation is recalled, otherwise forgotten. We simulated 1000 probability estimates for each combination of forgetting rate and outcome probability. Overall, qualitative misses were found to be extremely unlikely, only occurring (at a low rate) when forgetting rates were extremely high.

⁸ Equivalent results were observed when introducing random forgetting via a probability parameter f (see Costello & Watts, 2019) instead of estimating the weight λ on the prior. The reason for the similar results is that both λ and f serve the same role of modulating the weight of the prior probability relative to the observed sample and therefore the regressive effect toward $\frac{1}{2}$.

Table 2
Overview of the Data Corpus

Study	N	Effect	Qualitative Miss
1 Berndsen, McGarty, Van Der Pligt, and Spears (2001, E2)	22	IC ₁	✗
2 Bulli and Primi (2006)	158	IC ₁	✓
3 Eder et al. (2011, E1, online-prediction)	24	IC ₁	✓
4 Eder et al. (2011, E2, no-load)	33	IC ₁	✓
5 Eder et al. (2011, E2, load)	33	IC ₁	✓
6 Eder et al. (2011, E3, negative)	21	IC ₁	✓
7 Eder et al. (2011, E3, positive)	21	IC ₁	✓
8 Fiedler et al. (1993, no-valence-group)	15	IC ₁	✓
9 Fiedler et al. (1993, no-valence-person)	15	IC ₁	✗
10 Fiedler et al. (1993, valence-group)	15	IC ₁	✗
11 Fiedler et al. (1993, valence-person)	15	IC ₁	✓
12 Hamilton and Gifford (1976, E2)	70	IC ₁	✓
13 Klauer and Meiser (2000, E2)	20	IC ₁	✗
14 Madey and Chasteen (2004, younger)	22	IC ₁	✗
15 Meiser and Hewstone (2001, group)	20	IC ₁	✓
16 Meiser and Hewstone (2001, town)	20	IC ₁	✗
17 Meiser (2003, E1, impression)	33	IC ₁	✗
18 Meiser (2003, E1, memory)	33	IC ₁	✓
19 Meiser and Hewstone (2006, (a), town X)	35	IC ₁	✓
20 Meiser and Hewstone (2006, (a), town Y)	35	IC ₁	✓
21 Mutter (2000, distraction-older)	24	IC ₁	✓
22 Mutter (2000, distraction-younger)	24	IC ₁	✓
23 Primi and Agnoli (2002, E1)	253	IC ₁	✗
24 Primi and Agnoli (2002, E2)	106	IC ₁	✓
25 Rodríguez-Ferreiro and Barberia (2017)	240	IC ₁	✓
26 Spears et al. (1985)	141	IC ₁	✗
Subtotal (Qualitative Misses)			$\frac{17}{26}$ (65%)
27 Fleig et al. (2017, E1, negative)	56	IC _S	✗
28 Fleig et al. (2017, E1, positive)	56	IC _S	✓
29 Fleig et al. (2017, E2, base-rate, positive)	36	IC _S	✗
30 Fleig et al. (2017, E2, standard, positive)	39	IC _S	✓
31 Meiser (2003, E2, impression, negative)	38	IC _S	✗
32 Meiser (2003, E2, impression, positive)	38	IC _S	✗
33 Meiser (2003, E2, memory, negative)	38	IC _S	✗
34 Meiser (2003, E2, memory, positive)	38	IC _S	✗
35 Meiser and Hewstone (2004, E1, negative)	39	IC _S	✗
36 Meiser and Hewstone (2004, E1, positive)	39	IC _S	✓
37 Meiser and Hewstone (2004, E2, group, negative)	39	IC _S	✗
38 Meiser and Hewstone (2004, E2, group, positive)	39	IC _S	✗
39 Meiser and Hewstone (2004, E2, town, negative)	39	IC _S	✗
40 Meiser and Hewstone (2004, E2, town, positive)	39	IC _S	✓
41 Meiser and Hewstone (2004, E3, negative)	66	IC _S	✗
42 Meiser and Hewstone (2004, E3, positive)	66	IC _S	✗
43 Meiser et al. (2018, E1, positive)	66	IC _S	✓
44 Meiser et al. (2018, E2, visualization, negative)	30	IC _S	✗
45 Meiser et al. (2018, E2, visualization, positive)	30	IC _S	✓
46 Meiser et al. (2018, E3, pc, positive)	39	IC _S	✓
47 Meiser et al. (2018, E3, pc-salience, positive)	40	IC _S	✓
48 Meiser et al. (2018, E4, computer, positive)	64	IC _S	✓
49 Meiser et al. (2018, E4, self, positive)	63	IC _S	✓
Subtotal			$\frac{10}{23}$ (43%)
Total			$\frac{27}{49}$ (55%)

Note. Column “N” reports the number of participants in a given study. Column “Effect” indicates the specific illusory correlation being tested. Finally, column “Qualitative Miss” indicates whether participants estimated the majority outcome in the minority group as being the minority outcome in that group.

Figure 4. Model fits for the studies listed in Table 2. The ‘model failure’ quadrants correspond to regions in which there is a qualitative mismatch between model predictions and participants’ judgements.

investigation shows is that the optimality of the Rule of Succession and the prediction of illusory correlations only holds under very specific circumstances and conceptualizations

that are not well justified. The Rule of Succession often failed to account for Type-S illusory correlations. Indeed, we applied the rule to cases outside the domain of IC_1 effects for which it was originally proposed. Yet, it is not clear what kind of rationale would permit a demarcation between IC_1 and IC_S effects. Additionally, the rule also fails to account for the qualitative misses that are often found in the minority-group estimates. These misses show that illusory correlations involve something more than a regression to the $\frac{1}{2}$ -midpoint. Altogether, the theoretical and empirical issues discussed here raise the question whether a normative account of illusory correlations is possible. In the section below, we will answer this question affirmatively.

A Bayesian Marginal Account

As previously discussed, Fiedler and colleagues proposed that the observed illusory correlations are due to the use of a simple *pseudocontingency heuristic* that can be applied to marginal frequencies (for an overview, see Fiedler et al., 2013): *If two things occur often then assume that they are associated*. Because most instances come from Group A and most outcomes are ‘+’, Group A is thus assumed to be associated with the occurrence of ‘+’. According to Fiedler et al. (2009), this heuristic is broadly used and is claimed to underlie several phenomena, like the perception of ecological correlations (Hammond, 1973). The case for this heuristic typically revolves around its merits under the lens of ecological rationality: Even though one cannot *conclude* the presence of an association between the variables based on the marginal frequencies, this heuristic succeeds in capturing the sign of actual contingencies found in different environments (Kutzner et al., 2011).

The goal of this section is to show that one can make a case for this so-called heuristic within the tenets of Bayesian inference and to demonstrate its ability to overcome the limitations of the Rule of Succession. For this purpose, we will elaborate on an early proposal by Klauer (2015), which we will refer to as the *Bayesian Marginal Model*. We will explore the model’s viability vis-a-vis the corpus of data collected. The reason for using

Bayesian inference as a normative framework is that it provides us with a way to coherently represent beliefs regarding latent quantities such as population-level probabilities, and how these beliefs should be informed by incoming data (Griffiths & Tenenbaum, 2005; Knill & Richards, 1996; Oaksford & Chater, 1994; Tauber, Navarro, Perfors, & Steyvers, 2017). Importantly, note that the assumed reliance on Bayesian inference is predicated on the assumption that the reasoning processes used when generating probability estimates approximates the process of taking samples from a distribution (for an overview, see Sanborn & Chater, 2016). Also note that Bayesian accounts like the one developed below are rational in the sense that they lead to *coherent* sets of beliefs that cannot be exploited by an outsider through ‘Dutch-Book’ scenarios (Teller, 1973). This coherence is independent of whether they are optimal or not, which is something that will ultimately hinge on the match between the priors and the environment, among other aspects.

The model proposed explicitly adopts a subjectivist view of probability, representing the relative strength of (rational/coherent) beliefs in terms of probability *distributions*. This contrasts with the Rule of Succession as advocated by Costello and Watts (2019), who focused on point estimates and utilized their relative accuracy in known environments as a yardstick for rationality. These differences are reflected in our discussion of the Bayesian Marginal Model’s IC predictions, which focuses on the emergence of IC effects at the level of the posterior distributions of p_A and p_B . Thus, if participants base their probability estimates on a sample from the posterior distribution (and their frequency estimates on a sample of the posterior predictive distribution for the cell frequencies), then our predictions apply for the observed data without the need to postulate that people use a loss function of any kind or that groupwise proportions p in people’s environment follow a uniform distribution.

The model, its properties, and its performance

According to the pseudocontingency account (e.g., Fiedler et al., 2009), contingencies are inferred based on the association of skewed marginal frequencies. Even though marginal frequencies do not determine cell frequencies and thus the contingency, they do restrict the range of possible values (Duncan & Davis, 1953). Therefore, our model jointly considers the four cells of a 2×2 contingency table (see Table 1). It takes all possible cell distributions, with cell frequencies $\mathbf{k} = (k_{00}, k_{01}, k_{10}, k_{11})$ and cell probabilities $\mathbf{p} = (p_{00}, p_{01}, p_{10}, p_{11})$, into account that are in line with the observed marginal frequencies. These potential cell distributions and thus contingencies are weighted depending on the marginal frequencies and the reasoner's prior beliefs. The resulting output is a (posterior) belief about the contingency, or broadly speaking, the four cells will be reconstructed by the marginals.

Like any Bayesian account, our model can be broken down into three main components:

- A prior distribution $P(\mathbf{p})$ over the probability vector $\mathbf{p}$,
- a likelihood, here given by the multinomial probability mass function $P(\mathbf{k} | \mathbf{p})$, and
- a posterior distribution $P(\mathbf{p} | \mathbf{k})$ of the probability vector $\mathbf{p}$ given the data $\mathbf{k}$.

The prior distribution captures our beliefs regarding the probabilities $\mathbf{p}$ before encountering data $\mathbf{k}$. Although infinitely many priors are possible, we will focus our discussion on ‘uninformed’ priors that do not establish a predominance of any of the four table cells. Given the assumption that the data follow a multinomial distribution, we use a Dirichlet distribution to capture prior beliefs:

$$g(\mathbf{p}) = \frac{\Gamma(\sum \alpha_{ij})}{\prod \Gamma(\alpha_{ij})} \prod p_{ij}^{\alpha_{ij}-1}, \quad (6)$$

where $\boldsymbol{\alpha} = (\alpha_{00}, \alpha_{01}, \alpha_{10}, \alpha_{11})$, with all $\alpha_{ij} > 0$, is a vector of concentration parameters that can be interpreted as pseudocounts on each of the table cells. The expected prior p_{ij}

under this distribution is $\frac{\alpha_{ij}}{\sum \alpha_{ij}}$. Setting all α_{ij} to the same value for all i and j yields a prior distribution in which all joint outcomes are expected to be equiprobable. The larger this common value is, the stronger the prior beliefs will be (just as in the case of α and β parameters in the Beta distribution). The use of prior beliefs that are unbiased towards one group and/or outcome are reasonable, when studies use variables for which little to no prior information is given, as is the case of the studies included in our corpus. The motivation for using the Dirichlet distribution is twofold: First, it is an extremely flexible distribution, able to capture a multitude of beliefs via $\boldsymbol{\alpha}$. Second, while in the Rule of Succession estimates are computed per group, our model considers all four cell frequencies jointly and thus contingencies. Therefore, the generalization of the Beta distribution discussed earlier is needed. The Dirichlet distribution is the conjugate prior of the multinomial distribution. This means that when encountering multinomial-distributed data with probability mass function

$$P(\mathbf{k} \mid \mathbf{p}) = \binom{M}{k_{00}, k_{01}, k_{10}, k_{11}} \prod p_{ij}^{k_{ij}}, \quad (7)$$

the updated posterior beliefs $P(\mathbf{p} \mid \mathbf{k})$ will also be Dirichlet distributed, with parameter vector $\boldsymbol{\alpha}^* = \boldsymbol{\alpha} + \mathbf{k}$.

The prior and posterior Dirichlet distributions implicitly capture the reasoner's beliefs about the presence of an association between the variables (i.e., ϕ). The left panel of Figure 5 shows the case of $\boldsymbol{\alpha} = (1, 1, 1, 1)$, which yields a prior ϕ distribution whose mass is dispersed over the $[-1, 1]$ range but still more concentrated around zero. The center and right panels of Figure 5 illustrate the posterior ϕ distributions that are obtained when individuals update their beliefs based on the joint frequencies reported in Table 1. In the case of the middle panel (IC₁ example), we see that the posterior distribution is more peaked around $\phi = 0$, consistent with the fact that $P(+|A) = \frac{27}{36} = P(+|B) = \frac{9}{12}$. In the right panel (IC_Sexample), most of the posterior mass is on the negative range, consistent with the fact that $P(+|A) = \frac{16}{24} < P(+|B) = \frac{8}{8}$ and that the sample's ϕ is -0.33 .

Figure 5. Prior ϕ under $\alpha = (1, 1, 1, 1)$ and the subsequent posterior distributions when updating the prior with the joint frequencies reported in Table 1.

Now, let us turn to the updating of beliefs based on the observation of *marginal frequencies*. Although marginals are insufficient to determine the underlying joint frequencies, they nevertheless *constrain the range of possibilities*. To see this, note that given some set of marginal frequencies, the joint frequencies are completely determined if one of the joint frequencies is known. For instance, if we know k_{00} , then $k_{01} = k_{0\bullet} - k_{00}$, $k_{10} = k_{\bullet 0} - k_{00}$, and $k_{11} = M - k_{0\bullet} - k_{\bullet 0} + k_{00}$. It follows that we can generate all possible sets of joint frequencies by varying k_{00} between $\max(0, k_{0\bullet} + k_{\bullet 0} - M)$ and $\min(k_{0\bullet}, k_{\bullet 0})$. Given this relationship between joint and marginal frequencies, the probability of any set of marginal frequencies corresponds to:

$$P(k_{0\bullet}, k_{1\bullet}, k_{\bullet 0}, k_{\bullet 1} \mid \mathbf{p}) = \sum_{k_{00}=\max(0, k_{0\bullet}+k_{\bullet 0}-M)}^{\min(k_{0\bullet}, k_{\bullet 0})} P(k_{00}, k_{0\bullet}-k_{00}, k_{\bullet 0}-k_{00}, M-k_{0\bullet}-k_{\bullet 0}+k_{00} \mid \mathbf{p}), \quad (8)$$

where the probabilities on the right side are given by Equation 7.

Table 3 reports all possible joint frequencies under the marginals of the examples in Table 1. The associations (quantified by the ϕ -statistic) observed in most possible sets of joint frequencies turn out to be positive and the range of possible associations to be asymmetric, as associations can be as high as $\phi = 1$ but cannot go below $\phi = -0.33$. As

Table 3

Possible joint frequencies and associations (ϕ) under the marginals of the IC_1 and IC_S examples in Table 1

Type-1 Illusory Correlation Example (True $\phi = 0$)					
Index	k_{00}	k_{01}	k_{10}	k_{11}	ϕ
1	24	12	12	0	-0.33
2	25	11	11	1	-0.22
3	26	10	10	2	-0.11
4	27	9	9	3	0.00
5	28	8	8	4	0.11
6	29	7	7	5	0.22
7	30	6	6	6	0.33
8	31	5	5	7	0.44
9	32	4	4	8	0.56
10	33	3	3	9	0.67
11	34	2	2	10	0.78
12	35	1	1	11	0.89
13	36	0	0	12	1.00
Type-S Illusory Correlation Example (True $\phi = -0.33$)					
Index	k_{00}	k_{01}	k_{10}	k_{11}	ϕ
1	16	8	8	0	-0.33
2	17	7	7	1	-0.17
3	18	6	6	2	0.00
4	19	5	5	3	0.17
5	20	4	4	4	0.33
6	21	3	3	5	0.50
7	22	2	2	6	0.67
8	23	1	1	7	0.83
9	24	0	0	8	1.00

shown in Figure 6, this asymmetry in possible associations across the range of possible joint frequencies will be *greater* the more skewed the base rates are (see also Fiedler et al., 2013; Kareev, 1995).

Using Bayes' Theorem, we can combine Equations 6 and 8 to obtain the density of $\mathbf{p}$

Figure 6. Minimum, average, and maximum associations (quantified by the ϕ statistic) across different marginal frequencies ($k_{0\bullet} = k_{\bullet 0}$, with $M = 48$). When computing the average association for a given set of marginals, we attributed equal weights to all possible ϕ values.

conditional on the marginal frequencies and prior belief distribution:

$$\begin{aligned}
 g(\mathbf{p} \mid k_{0\bullet}, k_{1\bullet}, k_{\bullet 0}, k_{\bullet 1}, \boldsymbol{\alpha}) &\propto \sum_{k_{00}=\max(0, k_{0\bullet}+k_{\bullet 0}-M)}^{\min(k_{0\bullet}, k_{\bullet 0})} \binom{M}{k_{00}, k_{0\bullet}-k_{00}, k_{\bullet 0}-k_{00}, M-k_{0\bullet}-k_{\bullet 0}+k_{00}} \times \\
 &\quad \frac{\prod \Gamma(\alpha_{i,j} + k_{i,j})}{\Gamma(\sum(\alpha_{i,j}) + M)} \times g(\mathbf{p} \mid \boldsymbol{\alpha}^*(k_{00})).
 \end{aligned} \tag{9}$$

This posterior density is a mixture of the posterior Dirichlet distributions $g(\mathbf{p} \mid \boldsymbol{\alpha}^*)$ with $\boldsymbol{\alpha}^*(k_{00}) = \boldsymbol{\alpha} + \mathbf{k}'(k_{00})$, where $\mathbf{k}'(k_{00})$ corresponds to the vector of joint frequencies in the contingency table consistent with the given marginals and the given k_{00} . Each mixture weight (which is proportional to the first two multiplicative terms) is a function of the marginal frequencies and prior $\boldsymbol{\alpha}$. The second row of panels in Figure 7 illustrates the weights associated with the IC_1 example in Table 3 under different priors. The different joint-frequency sets are equally weighted when all $\alpha_{i,j} = 1$. Moving away from this prior

leads to differential weighting. For instance, when all $\alpha_{i,j} < 1$, the more ‘extreme’ joint frequencies are overweighted, with particular emphasis on the joint-frequencies that produce the largest ϕ . The reason behind this asymmetric overweighting is the fact low $\alpha_{i,j}$ values push most mass to extreme probabilities (Figure 1 provides an illustrative example, namely a Beta distribution with $\alpha = \beta = \frac{1}{2}$). Because the marginal frequencies are skewed, there is only one possible table that is consistent with these extreme probabilities, hence its overweighting.

Taking a single sample from this mixture density is straightforward: We first sample from a discrete distribution, with probabilities corresponding to the mixture weights. The sampled integer determines the component Dirichlet distribution from which to subsequently sample probabilities from. Finally, one can use these probabilities to sample a set of joint frequencies of size M . Multiple samples can be taken by repeating these steps.⁹ The third row of panels in Figure 7 illustrates the posterior ϕ distributions for the IC₁ example in Table 1 under different priors. We see that illusory correlations will be most prevalent in cases where the reasoner does not have strong prior beliefs against an association (e.g., when all $\alpha_{ij} = 1$). When such beliefs do in fact exist (e.g., when all $\alpha_{ij} = 50$), the propensity and magnitude of illusory correlations will be reduced. For example, consider a reasoner who learns that a) most students in a classroom perform well in tests, and b) that most of these students do their homework. Under the notion that an association is not implausible, the reasoner is likely to infer a positive association. In contrast, it seems somewhat implausible a priori that either of these observations are associated with students liking chocolate, which would justify a prior ϕ that is more strongly peaked on zero. Under such prior, it is unlikely that one would subsequently infer an association between student performance and liking chocolate, even after learning that

⁹ Note that the output of the Bayesian Marginal Model is a posterior distribution, not a point estimate. As the model is concerned with epistemic rationality as compared with optimality, it is not committed to the use of a specific loss function or any loss function at all. Also note, we focus our model evaluation on the ability to capture qualitative patterns rather than the relative precision of point predictions, wherefore our analysis is also not committed to some specific point estimator.

Figure 7. Illustration of the Bayesian Marginal Model and its predictions under different priors for the IC₁ example given in Table 1.

most students like chocolate. Moreover, note that the occurrence of qualitative misses is far from rare. Instead, its occurrence is quite robust across priors (see the bottom row of

Figure 7 depicting posterior probabilities for the minority group B). This follows from the fact that $\frac{k_{10}}{k_{1\bullet}} < \frac{1}{2}$ in 46% of the possible joint frequencies that are consistent with marginals found in the IC₁ example (see Table 3). A ‘problematic’ phenomenon under the Rule of Succession turns out to be an expected result under the Bayesian Marginal Model.

Given the potential of the Bayesian Marginal Model, we investigated its ability to capture the illusory correlations and qualitative misses observed in our study corpus. We restricted the α_{ij} priors to take on the same value, which means that *the model only has one free parameter*. This is analogous to the constraint imposed in the generalization of the Rule of Succession when fitting it to the same study corpus. In fact, our parameter restriction can be represented in the exact same way, with $\alpha_{ij} = \lambda \cdot \pi_{ij}$, with $\pi_{ij} = \frac{1}{4}$ for $i = 1, 2$ and $j = 1, 2$, where λ is a free parameter quantifying the concentration or peakedness of the prior.¹⁰

The results depicted in Figure 4 show that, overall, the Bayesian Marginal Model performs better than the Rule of Succession and its generalization in the sense that it can capture qualitative misses. Taken together, the assumption that individuals estimate probabilities based on marginal frequencies is able to account for the qualitative patterns found in the data. Put differently, one can make a rational as well as an empirical case for the pseudocontingency account proposed by Fiedler and colleagues.

Dependencies Across Probability Estimates, Skewed Beliefs, and New Data Formats

One of the main virtues of normative models is the fact that they often introduce new testable predictions that would have not been considered otherwise. We explore a number of such predictions concerning dependencies between the frequency estimates, the role of skewed beliefs, and new data formats.

¹⁰ For the sake of comparability, we fitted the Bayesian Marginal Model in exactly the same way as the Rule of Succession, using a constrained least-squares algorithm based on the model’s predictions for posterior means.

Dependencies. As previously stated, the samples taken under the Bayesian Marginal Model come from a mixture of Dirichlet distributions, each associated with one possible contingency table. Across these tables, the joint frequencies are negatively correlated as they sum to M , such that whenever k_{00} is small/large, k_{10} is likely to be large/small. As an example, we generated posterior samples for the IC₁ example in Table 1. We used a different α prior per sample (once again, with all α_{ij} being equal), which was taken from a mixture between a Gamma distribution with shape and rate parameters 1.26 and 1.16, respectively, and a uniform distribution between 0 and 50. The mixture weights were .63 and .37, respectively. This mixture distribution was found to be the one best fitting the α estimates obtained with the study corpus. We found the resulting $\hat{p}_A$ and $\hat{p}_B$ samples to be negatively correlated (Spearman’s $r = -0.67$).

In contrast, the probability estimates $\hat{p}$ coming from the Rule of Succession are assumed to be computed separately for each group, with no presumption of any kind of crosstalk. This means that $\hat{p}_A$ and $\hat{p}_B$ should be *stochastically independent* (i.e., the expected correlation is 0). We tested these differing predictions by fitting a regression model to the IC_S data coming from Klauer and Meiser (2000) and Meiser et al. (2018, see Table 2).¹¹ In addition to having $\hat{p}_A$ as a predictor of $\hat{p}_B$, we also introduced the predictor ‘Experiment’. The estimated coefficient was $\beta_{\hat{p}_A} = -0.25$ ($SE = 0.08$, $p = .001$), supporting the predictions made by the Bayesian Marginal Model.¹²

Skewed Beliefs. So far, we assumed that all parameter values of prior α are equal, which is reasonable when participants know nothing about the relative frequency of groups and/or outcomes beforehand. However, when skewed frequencies are expected *a priori*, illusory correlations are only expected under certain conditions. Figure 8 shows the

¹¹ We thank Thorsten Meiser for making the data from Meiser et al. (2018) available.

¹² As a caveat, note that the prediction assumes (a) equal $\alpha_{i,j}$ and (b) that respondents do not differ in their propensities or response biases to produce generally high or low frequency estimates. The first assumption is plausible where groups and outcomes are abstract or novel. The second assumption is more plausible when looking at intraindividual correlations across several IC tasks administered to the same person.

posterior ϕ distributions under different priors that respect the 3:1 ratio found in the IC₁ example in Table 1. The figure contains examples involving skewed outcome priors with outcome ‘+’ cells being overweighted relative to outcome ‘−’ cells (e.g., $\alpha = (3, 1, 3, 1)$) and skewed group-membership priors (e.g., $\alpha = (3, 3, 1, 1)$). As can be seen, the positive bias in the posterior ϕ distribution, indicating an IC₁ effect, vanishes the more informative the α becomes. Consequently, illusory correlations are expected to be reduced or even absent whenever individuals have clear prior beliefs regarding the actual skewness of the event distributions that is to be experienced. Similarly, no IC₁ effect is expected, when prior beliefs correspond to the actual frequencies (e.g., $\alpha = (9, 3, 3, 1)$). Finally, Figure 8 also depicts posterior ϕ distributions under priors that overweight the ‘wrong’ cells (e.g., $\alpha = (1, 1, 3, 3)$ and $\alpha = (1, 3, 1, 3)$) which are expected to increase illusory correlations. Therefore, future work should consider systematically investigating the role of prior beliefs in the occurrence of illusory correlations.

Asymmetry. Asymmetric illusory correlation effects are documented in the literature, with IC effects tending to be greater in contexts that are mostly comprised of positive outcomes (e.g., Fleig et al., 2017; Meiser et al., 2018). These asymmetries are observed in experimental designs in which the overall positivity/negativity is manipulated within participants. In such scenarios, there are two contexts (e.g., Town X and Town Y). The marginal frequencies of groups and outcomes are skewed within each context, and co-vary across contexts: what is frequent within one context is infrequent within the other context and vice versa. One possible explanation for this asymmetry is that the overall positivity/negativity of each context is detected quickly by the participants, which leads them to subsequently make inferences about each group in each context using skewed priors: In both contexts, the most frequent outcome will be overweighted relative to the cells corresponding to the infrequent outcome. In terms of the Bayesian Marginal Model, such priors would reduce the magnitude of the IC_S effects (see the first row of Figure 8). Given people’s tendency to attend more to losses than gains (Yechiam & Hochman, 2013),

Figure 8. Posterior ϕ for the IC₁ example in Table 1 according to the Bayesian Marginal Model under different skewed priors.

one would expect a greater skew in the context where negative outcomes are most frequent, which in turn would produce the observed asymmetry.

New Response Formats. Given that the output of the Bayesian Marginal Model is a posterior distribution, the model readily generates predictions for alternative response formats for which the Rule of Succession, relying on one particular point estimate, remains silent. For example, participants should not only be able to generate point estimates for the

probabilities $P(+|A)$ and $P(+|B)$, they should also be able to estimate the probabilities that each probability falls into given ranges, and these range estimates should correspond to the model's predictions derived from the posterior distribution. Again, algorithmically, the range estimate could be based on mental sampling from the posterior distribution.

New Data Formats. A final prediction illustrates that the model is readily extended to account for findings using new data formats. For instance, the model can be adapted to the case where data from just one cell of the contingency table is presented. Consider a scenario in which participants learn that 21 of 49 cases under study are desirable behaviors of Group A members. This information necessitates neither skewed marginals nor a contingency. Yet, under flat priors, it should bias posterior beliefs towards an illusory correlation.

Taken together, the Bayesian Marginal Model does not only show explanatory sufficiency (it accounts for the major qualitative patterns in the data corpus) and a normative backing, it is also a fruitful theory, readily generating new predictions, including the effects of prior beliefs (see Coenen et al., 2018).

General Discussion

Theoretical work on illusory correlations has relied on one of two main interpretations. On the one hand, illusory correlations have been argued to reflect biased or at least normatively unjustified information processing. On the other hand, Costello and Watts (2019) argued that illusory correlations reflect optimal, mathematically correct and hence rational responses. The present work raises a number of issues that ultimately speak against both theoretical interpretations: First, a sensible account of a specific illusory correlation (e.g., IC_1) should also be able to accommodate closely-related variants (e.g., IC_S). Second, any successful theoretical account needs to go beyond a single effect and consider other qualitative patterns in the data. By focusing on illusory correlation effects in isolation, one is ignoring the fact that the data often show qualitative misses that cannot

be accommodated by models that assume some kind of regression towards the mean/prior. Third, we nevertheless disagree that illusory correlations reflect an irrational, normatively unjustified kind of information processing: The Bayesian Marginal Model provides a normative reconstruction of the pseudcontingency heuristic, predicting IC_1 and IC_S effects based on the use of marginal frequencies following the norms of Bayesian belief updating. Furthermore, this model turned out to be able to account for the qualitative patterns in the study corpus and successfully predict the presence of dependencies across individual probability estimates. The results demonstrate that a successful normative account of illusory correlations is possible.

The occurrence of illusory correlations in Figure 7 should be sufficient to dispel the notion that the pseudocontingency heuristic proposed by Fiedler and colleagues (e.g., Fiedler et al., 2009) is logically invalid. Marginal frequencies do carry valid information about the joint frequencies, otherwise, in the Bayesian Marginal Model, the posterior beliefs would match the priors. Under weakly informative priors, skewed marginal frequencies are expected to result (via Bayes' theorem) in the belief that both variables are associated.

Between the two contrary interpretations of illusory correlations as irrational and unjustified versus 'mathematically correct' and rational, we therefore argue for a third possibility, namely that they reflect an instance of *bounded rationality* (Simon, 1990). According to this view, limitations of the human information-processing system introduce constraints that (a) prevent human reasoners and decision makers from reaching overall optimal solutions, but (b) do not prevent them from finding normatively adequate solutions within the smaller search space of solutions consistent with these constraints. In the present case, an optimally tuned information-processing system would make full use of the presented joint frequencies instead of only the marginal frequencies. A human decision maker's memory limitations, however, constrain the number of cases in which the realization of two variables in an event are bound together (e.g., Bays, Wu, & Husain, 2011; Treisman, 1998; Wheeler & Treisman, 2002), which might discourage individuals to

use them in inferential judgments. Eder et al. (2011) have reported results supporting this hypothesis, in which they show that the impairment of working memory exacerbates the occurrence of illusory correlations. In other cases, a reasoner might often encounter the different variables separately, sometimes even at different points in time. Given these limitations, the Bayesian Marginal Model holds that boundedly rational reasoners fall back on the marginal frequencies from which illusory correlations are predicted as an outcome of a normatively adequate, Bayesian use of the information in these marginal frequencies.

According to Marr’s (1982) classification, the Bayesian Marginal Model provides a characterization at the *computational level*. It is assumed that individuals have the capacity to perform or mimic the normatively appropriate computations and use the normative prescriptions to guide their choices and behaviors. At no point does it make any claims at the *algorithmic level*, that is, claims regarding the mechanisms by which the aforementioned computations are performed or mimicked (for a relevant discussion, see Van Rooij, 2008). We find modeling at the computational level to be the most appropriate here given the nature of the data available: In illusory-correlation experiments, each participant contributes very few judgments. This means that one does not have sufficient degrees of freedom to test theoretical accounts that make fine-grained algorithm-level claims. To make matters worse, aggregating data across participants is very likely to introduce distortions that can spuriously reject models and lead researchers astray (for recent discussions, see Kellen & Klauer, 2019; Regenwetter & Robinson, 2017).

Addressing Concerns on Model Flexibility, Evaluation, and Comparison

The minimal requirement for any candidate model is that it succeeds in capturing the data at hand (Bogen & Woodward, 1988). Given that we are dealing with aggregate data, it is reasonable to focus our evaluation on the models’ ability to capture the qualitative patterns found in the data (for discussions, see Kellen, 2019; Navarro, 2019; Shiffrin & Nobel, 1997). It turns out that the Rule of Succession (as well as its generalization) cannot

fulfill this requirement. This reason alone is why the Rule of Succession has to be dismissed; not because our comparison between the Rule of Succession and the Bayesian Marginal Model might place the former model at a disadvantage due to not taking potential differences in flexibility into account (e.g., Kellen, Klauer, & Bröder, 2013; Myung & Pitt, 2018). Note also that in our fits of the Rule of Succession and the Bayesian Marginal Model to the data corpus, we used only one free parameter, governing the peakedness of the prior.

The Bayesian Marginal Model succeeded where the Rule of Succession failed. However, Bayesian models are notorious for their extreme flexibility, due to the possibility of establishing priors and likelihoods in so many different, arbitrary ways. This may render their eventual success as nothing short of certain (Bowers & Davis, 2012; see also Coenen, Nelson, & Gureckis, 2019). We argue that the development history of the Bayesian Marginal Model successfully addresses this concern. First, the model is a formal instantiation of the pseudocontingency heuristic proposed much earlier by Fiedler and colleagues (e.g., Fiedler et al., 2009, 2013). This means that the likelihood function of the model was constrained from the onset. As sketched out by Klauer (2015), the model is able to yield IC_1 effects under non-informative priors, introducing the notion that the occurrence of illusory correlations does not require any kind of bias in people’s prior beliefs. In the present paper, we went beyond Klauer’s first proof of concept and tested the model against a large corpus of published studies while keeping the basic components of the model fixed. The only part of the model that was ‘free to vary’ across studies was a *single* concentration parameter λ , which determines the strength of the prior (here non-informative beliefs, i.e., prior beliefs that are kept unbiased towards any specific group or outcome).¹³ Moreover, the Bayesian Marginal Model *successfully predicted* the presence of (a) qualitative misses and (b) negative dependencies across probability estimates. The model is fruitful in generating a host of new predictions, as illustrated in the section

¹³ Note that the use of non-informative priors was accompanied by a focus on studies in which such priors are deemed highly plausible. As previously discussed, the studies included in the corpus involved variables for which little to no prior information is available.

“Dependencies Across Probability Estimates, Skewed Beliefs, and New Data Formats”.

Decisions from Experience and Skewed Sampling

The learning phase in illusory-correlation experiments shares important similarities with some of the paradigms used in the study of *experience-based decisions* (e.g., Hertwig et al., 2004; Kellen, Pachur, & Hertwig, 2016). As an example, consider the following options:

$$\mathcal{A} = \left(\begin{array}{cc} \$120 & \$0 \\ .70 & .30 \end{array} \right) \quad \mathcal{B} = \left(\begin{array}{cc} \$100 & \\ 1 & \end{array} \right)$$

Option $\mathcal{A}$ is a lottery stating the yields \$120 with probability .70, otherwise nothing.

Option $\mathcal{B}$ always yields \$100. Note that $\mathcal{B}$ has a greater expected value than $\mathcal{A}$ (\$100 and \$84, respectively). In a typical *sampling paradigm*, both options’ outcomes and outcome probabilities are learned by sampling random draws from each of them before the participant makes a final, consequential choice (Hertwig et al., 2004).

If option $\mathcal{A}$ is sampled more often than option $\mathcal{B}$, one might infer an association between the majorly-sampled option and monetary gains that would increase $\mathcal{A}$ ’s relative attractiveness. Note that a preference for $\mathcal{A}$ over $\mathcal{B}$ would deviate from people’s choices when options are described (e.g., Hertwig et al., 2004). The potential effects of skewed sampling in experience-based choices has received minimal attention by researchers so far (e.g., Bott & Meiser, in press).¹⁴ This situation is partly due to participants being allowed to freely sample both options as often and in any order they desire. However, this freedom ignores the fact that in realistic settings, some options are experienced more often than others, without any control on the decision-maker’s part.

¹⁴ Mostly, investigated effects referred to how representative or veridical the options’ sampled probabilities are (e.g., Hau, Pleskac, Kiefer, & Hertwig, 2008; Hilbig & Glöckner, 2011), rather than testing the effect of sample sizes per option on preferences.

Final Thoughts

Finally, let us highlight the fact that the Bayesian Marginal Model proposed here attributes a passive observer role to the reasoner. This situation implies that many phenomena attributed to the actions of the reasoner in a given environment, such as the way in which different options are actively sampled, or to the environment explored, are outside its current scope (for an overview on the topic of information search, see Todd, Hills, & Robbins, 2012). For instance, Le Mens and Denrell (2011) considered the case in which two options are objectively equally attractive, but asymmetric with regard to information: We can learn about one of the options only when choosing it (*selective access*), whereas we can learn about the other irrespective of whether it is chosen or not (*systematic access*). This asymmetry will lead to a preference for the latter option, even if the decision maker is rational and knows the structure of the environment.

Likewise in naturalistic settings, different groups differ in how informational access is controlled. For some, we are bombarded with relevant samples in the news media almost everyday; for others, we have to make an effort to gather the relevant bits of information. Such informational asymmetries and active sampling biases (e.g., Fiedler & Wänke, 2009) likely interact with the formation of illusory correlations. The integration of different Bayesian accounts thereof into a single, comprehensive account constitutes an important theoretical challenge that is yet to be tackled.

References

- Aydogan, I. (2019). *The role of prior beliefs and their updating in decisions under experienced ambiguity*. Manuscript submitted for publication.
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49(6), 1622–1631.
- Berndsen, M., McGarty, C., Van Der Pligt, J., & Spears, R. (2001). Meaning-seeking in the illusory correlation paradigm: The active role of participants in the categorization process. *British Journal of Social Psychology*, 40(2), 209–233. doi: 10.1348/014466601164821
- Bogen, J., & Woodward, J. (1988). Saving the phenomena. *The Philosophical Review*, 97(3), 303–352.
- Bott, F. M., & Meiser, T. (in press). Pseudocontingency inference and choice: The role of information sampling. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in Psychology and Neuroscience. *Psychological Bulletin*, 138(3), 389–414. doi: 10.1037/a0026450
- Buchak, L. (2013). *Risk and rationality*. Oxford, UK: Oxford University Press.
- Bulli, F., & Primi, C. (2006). Illusory correlation and cognitive processes: A multinomial model of source-monitoring. *Review of Psychology*, 13(2), 95–102.
- Busemeyer, J. R., & Bruza, P. D. (2012). *Quantum models of cognition and decision*. Cambridge, UK: Cambridge University Press.
- Carnap, R. (1952). *The continuum of inductive methods*. Chicago, IL: Chicago University Press.
- Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, 74(3), 271–280. doi: 10.1037/h0027592
- Cheng, P. W., & Novick, L. R. (1990). A probabilistic contrast model of causal induction.

- Journal of Personality and Social Psychology*, 58(4), 545–567.
- Coenen, A., Nelson, J. D., & Gureckis, T. M. (2019). Asking the right questions about the psychology of human inquiry: Nine open challenges. *Psychonomic Bulletin and Review*, 26, 1548–1587. doi: 10.3758/s13423-018-1470-5
- Costello, F., & Watts, P. (2014). Surprisingly rational: Probability theory plus noise explains biases in judgment. *Psychological Review*, 121(3), 463–480. doi: 10.1037/a0037010
- Costello, F., & Watts, P. (2016). People’s conditional probability judgments follow probability theory (plus noise). *Cognitive Psychology*, 89, 106–133. doi: 10.1016/j.cogpsych.2016.06.006
- Costello, F., & Watts, P. (2019). The rationality of illusory correlation. *Psychological Review*, 126(3), 437–450. doi: 10.1037/rev0000130
- DeGroot, M. H. (1975). *Probability and statistics*. Reading, Mass.: Addison-Wesley.
- Dougherty, M. R., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review*, 106(1), 180–209. doi: 10.1037/0033-295X.106.1.180
- Duncan, O. D., & Davis, B. (1953). An alternative to ecological correlation. *American Sociological Review*, 18(6), 665–666.
- Dzhafarov, E. N., & Kujala, J. V. (2016). Context–content systems of random variables: The Contextuality-by-Default theory. *Journal of Mathematical Psychology*, 74, 11–33. doi: 10.1016/j.jmp.2016.04.010
- Eder, A. B., Fiedler, K., & Hamm-Eder, S. (2011). Illusory correlations revisited: The role of pseudocontingencies and working-memory capacity. *Quarterly Journal of Experimental Psychology*, 64(3), 517–532. doi: 10.1080/17470218.2010.509917
- Elqayam, S., & Evans, J. S. B. T. (2011). Subtracting ought from is: Descriptivism versus normativism in the study of human thinking. *Behavioral and Brain Sciences*, 34(5), 233–248. doi: 10.1017/S0140525X1100001X

- Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review*, *103*(1), 193–214. doi: 10.1037/0033-295X.103.1.193
- Fiedler, K. (2000). Illusory correlations: A simple associative algorithm provides a convergent account of seemingly divergent paradigms. *Review of General Psychology*, *4*(1), 25–58. doi: 10.1037/1089-2680.4.1.25
- Fiedler, K. (2010). Pseudocontingencies can override genuine contingencies between multiple cues. *Psychonomic Bulletin & Review*, *17*(4), 504–509. doi: 10.3758/PBR.17.4.504
- Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology*, *87*(4), 453–467. doi: 10.1037/0022-3514.87.4.453
- Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: An integrative account of an intriguing cognitive illusion. *Psychological Review*, *116*(1), 187–206. doi: 10.1037/a0014480
- Fiedler, K., Kutzner, F., & Vogel, T. (2013). Pseudocontingencies: Logically unwarranted but smart inferences. *Current Directions in Psychological Science*, *22*(4), 324–329. doi: 10.1177/0963721413480171
- Fiedler, K., Russer, S., & Gramm, K. (1993). Illusory correlations and memory performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 111–136.
- Fiedler, K., & Wänke, M. (2009). The cognitive-ecological approach to rationality in social psychology. *Social Cognition*, *27*(5), 699–732.
- Fleig, H., Meiser, T., Ettlin, F., & Rummel, J. (2017). Statistical numeracy as a moderator of (pseudo)contingency effects on decision behavior. *Acta Psychologica*, *174*, 68–79. doi: 10.1016/j.actpsy.2017.01.002
- Gelman, A., & Carlin, J. (2014). Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors. *Perspectives on Psychological Science*, *9*(6), 641–651.

- doi: 10.1177/1745691614551642
- Gigerenzer, G. (2019). Axiomatic rationality and ecological rationality. *Synthese*. doi: 10.1007/s11229-019-02296-5
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90. doi: 10.1037/0033-295x.109.1.75
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive psychology*, 51(4), 334–384.
- Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology*, 12(4), 392–407. doi: 10.1016/S0022-1031(76)80006-6
- Hamilton, D. L., & Rose, T. L. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology*, 39(5), 832–845. doi: 10.1037/0022-3514.39.5.832
- Hammond, J. L. (1973). Two sources of error in ecological correlations. *American Sociological Review*, 38, 764–777.
- Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, 21(5), 493–518. doi: 10.1002/bdm.598
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539. doi: 10.1111/j.0956-7976.2004.00715.x
- Hilbert, M. (2012). Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making. *Psychological Bulletin*, 138(2), 211–237. doi: 10.1037/a0025940
- Hilbig, B. E., & Glöckner, A. (2011). Yes, they can! Appropriate weighting of small probabilities as a function of information acquisition. *Acta Psychologica*, 138(3),

- 390–396. doi: 10.1016/j.actpsy.2011.09.005
- Jeffreys, H. (1939). *Theory of Probability* (1st ed.). Oxford: The Clarendon Press.
- Johnson-Laird, P. N., Khemlani, S. S., & Goodwin, G. P. (2015). Logic, probability, and human reasoning. *Trends in Cognitive Sciences*, 19(4), 201–214. doi: 10.1016/j.tics.2015.02.006
- Kareev, Y. (1995). Positive bias in the perception of covariation. *Psychological Review*, 102(3), 490–502. doi: 10.1037/0033-295X.102.3.490
- Kellen, D. (2019). A model hierarchy for psychological science. *Computational Brain & Behavior*, 2(3), 160–165. doi: 10.1007/s42113-019-00037-y
- Kellen, D., & Klauer, K. C. (2019). Theories of the wason selection task: A critical assessment of boundaries and benchmarks. *Computational Brain & Behavior*, 1–13.
- Kellen, D., Klauer, K. C., & Bröder, A. (2013). Recognition memory models and binary-response ROCs: A comparison by minimum description length. *Psychonomic Bulletin and Review*, 20, 693–719. doi: 10.3758/s13423-013-0407-2
- Kellen, D., Pachur, T., & Hertwig, R. (2016). How (in) variant are subjective representations of described and experienced risk and rewards? *Cognition*, 157, 126–138.
- Kerman, J. (2011). *A closed-form approximation for the median of the beta distribution*. Retrieved from <https://arxiv.org/abs/1111.0433v1>
- Keynes, J. M. (1921). *A treatise on probability*. Dover Publications.
- Klauer, K. C. (2015). Mathematical modeling. In B. Gawronski & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (pp. 371–389). New York, NY, US: Guilford Press.
- Klauer, K. C., & Meiser, T. (2000). A source-monitoring analysis of illusory correlations. *Personality and Social Psychology Bulletin*, 26(9), 1074–1093. doi: 10.1177/01461672002611005
- Knill, D. C., & Richards, W. (1996). *Perception as bayesian inference*. Cambridge, Mass:

MIT Press.

- Kutzner, F., Vogel, T., Freytag, P., & Fiedler, K. (2011). Contingency inferences driven by base rates: Valid by sampling. *Judgment and Decision Making*, 6(3), 211–221.
- Laplace, P. S. (1820/1951). *Philosophical Essays on Probabilities*. New York: Dover.
- Le Mens, G., & Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological review*, 118(2), 379–392.
- Madey, S. F., & Chasteen, A. L. (2004). Age-related health stereotypes and illusory correlation. *International Journal of Aging and Human Development*, 58(2), 109–126. doi: 10.2190/81XB-QPYN-9ADU-5L11
- Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994). Illusory correlation in the perception of groups: An extension of the distinctiveness-based account. *Journal of Personality and Social Psychology*, 67(3), 414–429. doi: 10.1037/0022-3514.67.3.414
- Meiser, T. (2003). Effects of processing strategy on episodic memory and contingency learning in group stereotype formation. *Social Cognition*, 21(2), 121–156. doi: 10.1521/soco.21.2.121.21318
- Meiser, T. (2006). Contingency learning and biased group impressions. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition*. (pp. 183–209). Cambridge University Press.
- Meiser, T., & Hewstone, M. (2001). Crossed categorization effects on the formation of illusory correlations. *European Journal of Social Psychology*, 31(4), 443–466. doi: 10.1002/ejsp.55
- Meiser, T., & Hewstone, M. (2004). Cognitive processes in stereotype formation: The role of correct contingency learning for biased group judgments. *Journal of Personality and Social Psychology*, 87(5), 599–614. doi: 10.1037/0022-3514.87.5.599
- Meiser, T., & Hewstone, M. (2006). Illusory and spurious correlations: Distinct

- phenomena or joint outcomes of exemplar-based category learning? *European Journal of Social Psychology*, 36(3), 315–336. doi: 10.1002/ejsp.304
- Meiser, T., Rummel, J., & Fleig, H. (2018). Pseudocontingencies and choice behavior in probabilistic environments with context-dependent outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(1), 50–67. doi: 10.1037/xlm0000432
- Mutter, S. A. (2000). Illusory correlation and group impression formation in young and older adults. *Journals of Gerontology - Series B Psychological Sciences and Social Sciences*, 55(4), 224–237. doi: 10.1093/geronb/55.4.P224
- Myung, J. I., & Pitt, M. A. (2018). Model comparison in psychology. In E.-J. Wagenmakers (Ed.), *Stevens' handbook of experimental psychology and cognitive neuroscience. Volume 5: Methodology* (4th ed., p. 85-118). Hoboken, New Jersey: Wiley. doi: 10.1002/9781119170174.epcn503
- Navarro, D. J. (2019). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior*, 2(1), 28–34.
- 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. doi: 10.1037/a0017351
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608-631.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning* (1st ed.). Oxford [u.a.]: Oxford University Press.
- Primi, C., & Agnoli, F. (2002). Children correlate infrequent behaviors with minority groups: A case of illusory correlation. *Cognitive Development*, 17(1), 1105–1131. doi: 10.1016/S0885-2014(02)00076-X
- Regenwetter, M., & Robinson, M. M. (2017). The construct–behavior gap in behavioral

- decision research: A challenge beyond replicability. *Psychological Review*, 124(5), 533-550. doi: 10.1037/rev0000067
- Rodríguez-Ferreiro, J., & Barberia, I. (2017). The moral foundations of illusory correlation. *PLoS ONE*, 12(10), 1–10. doi: 10.1371/journal.pone.0185758
- Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in cognitive sciences*, 20(12), 883–893.
- Savage, L. J. (1954). *The foundations of statistics*. New York: Wiley.
- Sherman, J. W., Kruschke, J. K., Sherman, S. J., Percy, E. J., Petrocelli, J. V., & Conrey, F. R. (2009). Attentional processes in stereotype formation: A common model for category accentuation and illusory correlation. *Journal of Personality and Social Psychology*, 96(2), 305–323. doi: 10.1037/a0013778
- Shiffrin, R. M., & Nobel, P. A. (1997). The art of model development and testing. *Behavior Research Methods, Instruments, & Computers*, 29(1), 6–14.
- Simon, H. A. (1990). Bounded rationality. In J. Eatwell, M. Milgate, & P. Newman (Eds.), *Utility and probability* (pp. 15–18). London: Palgrave Macmillan UK. Retrieved from https://doi.org/10.1007/978-1-349-20568-4_5 doi: 10.1007/978-1-349-20568-4_5
- Skovgaard-Olsen, N., Kellen, D., Hahn, U., & Klauer, K. C. (2019). Norm conflicts and conditionals. *Psychological Review*, 126(5), 611-633. doi: 10.1037/rev0000150
- Spears, R., van der Pligt, J., & Eiser, J. R. (1985). Illusory correlation in the perception of group attitudes. *Journal of Personality and Social Psychology*, 48(4), 863–875. doi: 10.1037/0022-3514.48.4.863
- Spohn, W. (2012). *The laws of beliefs*. Oxford, UK: Oxford University Press.
- Tauber, S., Navarro, D. J., Perfors, A., & Steyvers, M. (2017). Bayesian models of cognition revisited: Setting optimality aside and letting data drive psychological theory. *Psychological review*, 124(4), 410-441.
- Teller, P. (1973). Conditionalization and observation. *Synthese*, 26(2), 218–258.

- Todd, P. M., Hills, T. T., & Robbins, T. W. (2012). *Cognitive search: Evolution, algorithms, and the brain*. Boston, Mass.: MIT press.
- Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 353(1373), 1295–1306.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131. doi: 10.1126/science.185.4157.1124
- Vadillo, M. A., Blanco, F., Yarritu, I., & Matute, H. (2016). Single-and dual-process models of biased contingency detection. *Experimental Psychology*, 63(1), 3–19. doi: 10.1027/1618-3169/a000309
- Van Rooy, D., Vanhooissen, T., & Van Overwalle, F. (2013). Illusory correlation, group size and memory. *Journal of Experimental Social Psychology*, 49(6), 1159–1167. doi: 10.1016/j.jesp.2013.05.006
- Van Rooij, I. (2008). The tractable cognition thesis. *Cognitive science*, 32(6), 939–984.
- Viscusi, W. K. (1989). Prospective reference theory: Toward an explanation of the paradoxes. *Journal of Risk and Uncertainty*, 2, 235–264. doi: 10.1007/BF00209389
- Vogel, T., Kutzner, F., Fiedler, K., & Freytag, P. (2013). How majority members become associated with rare attributes: Ecological correlations in stereotype formation. *Social Cognition*, 31(4), 427–442. doi: 10.1521/soco_2012_1002
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior (commemorative edition)*. Princeton, NJ: Princeton University Press.
- Wakker, P. P. (2002). Decision-principles to justify Carnap's updating method and to suggest corrections of probability judgments. *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence*, 544–551.
- Wasserman, E. A., Dorner, W. W., & Kao, S. F. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16(3), 509–521. doi:

10.1037/0278-7393.16.3.509

Wasserman, L. (2004). *All of statistics*. New York: Springer.

Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131(1), 48-64.

Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: Review and analysis of the unique effects of losses over gains. *Psychological Bulletin*, 139, 497-518.