SONDERFORSCHUNGSBEREICH 504

Rationalitätskonzepte,
Entscheidungsverhalten und
ökonomische Modellierung

No. 08-36

Pseudocontingencies: An Integrative Account of an Intriguing Cognitive Illusion

Klaus Fiedler*
and Peter Freytag**
and Thorsten Meiser***

January 2009

Authors Note. Klaus Fiedler and Peter Freytag, University of Heidelberg, Department of Psychology, Hauptstrasse 47, D-69117 Heidelberg, Germany. Thorsten Meiser, University of Marburg, Department of Psychology, Gutenbergstrasse 18, D-35032 Marburg, Germany. The present research was supported by several grants of the Deutsche Forschungsgemeinschaft to all three authors. Insightful comments by Jürgen Baumert, Nick Chater, Robyn Dawes, Yaakov Kareev, Joachim Krueger, Florian Kutzner, and Tobias Vogel are gratefully acknowledged. Correspondence concerning this manuscript should be addressed via email to: kf@psychologie.uni-heidelberg.de.

*Sonderforschungsbereich 504/ Universität Heidelberg, email: Klaus_Fiedler@psi-sv2.psi.uni-heidelberg.de

**Universität Heidelberg, email:

***Universität Marburg, email:

Universität Mannheim
L 13,15
68131 Mannheim

Pseudocontingencies: An Integrative Account of an Intriguing Cognitive Illusion

Klaus Fiedler
University of Heidelberg

Peter Freytag
University of Heidelberg

Thorsten Meiser
University of Marburg

Running Head: Pseudocontingencies

Authors' Note. Klaus Fiedler and Peter Freytag, University of Heidelberg, Department of Psychology, Hauptstrasse 47, D-69117 Heidelberg, Germany. Thorsten Meiser, University of Marburg, Department of Psychology, Gutenbergstrasse 18, D-35032 Marburg, Germany. The present research was supported by several grants of the Deutsche Forschungsgemeinschaft to all three authors. Insightful comments by Jürgen Baumert, Nick Chater, Robyn Dawes, Yaakov Kareev, Joachim Krueger, Florian Kutzner, and Tobias Vogel are gratefully acknowledged. Correspondence concerning this manuscript should be addressed via email to: kf@psychologie.uni-heidelberg.de.

Abstract

The term pseudocontingency (PC) denotes the logically unwarranted inference of a contingency between two variables X and Y from information other than pairs of x_i, y_i observations, namely, the variables' univariate baserates as assessed in one or more ecological contexts. We summarize recent experimental evidence, showing that PCs can play a pivotal role in many areas of judgment and decision making. We argue that the exploitation of the informational value of baserates underlying PCs offers an alternative perspective on many phenomena in the realm of adaptive cognition that have been studied in isolation so far. Although PCs can lead to serious biases under some conditions, they afford an efficient strategy for inductive inference making in probabilistic environments that render baserate information, rather than genuine covariation information, readily available.

Pseudocontingencies: An Integrative Account of an Intriguing Cognitive Illusion

Adaptive behavior hinges on our ability to extract meaningful regularities from the flood of information provided by the physical and social environment. In an uncertain world, discerning regularities often amounts to assessing statistical contingencies between two variables over a series of events, such as correlations between causes and events, signals and significant events, predictions and feedback, actions and consequences, or behaviors associated with particular personality types. It is no exaggeration to claim that the ability to assess contingencies is crucial for adaptive learning and behavior, for rational action and decision making, and – ultimately – for survival in a risky and uncertain world. Thus, contingency assessment is commonly considered a major module of inductive intelligence, as stated in several pertinent reviews (Allan, 1993; Allan, Hannah, Crump & Siegel, 2008; Alloy & Tabachnik, 1984; Arieh & Algom, 2002; Crocker, 1981; Fiedler, 2000).

In this article, we present an alternative perspective on inductive learning, drawing on pseudocontingencies rather than contingencies proper. The major claim we will put forward is that when organisms figure out correlations or contingencies (in case of nominally scaled variables), they often engage in a cognitive inference process that is sensitive to something different from a genuine correlation. To start with a definition, the term "pseudocontingency" denotes the (logically unwarranted) inference of a contingency between two variables X and Y from information other than pairs of x_i, y_i observations.

In empirical reality, there are many situations in which joint observations of two or more variables are not available, or in which environmental or mnemonic constraints preclude the use of such genuine contingency information. In these situations, pseudocontingencies (PCs) can be inferred from other information, particularly, from the covariation of unequal baserates of high and low levels of X and Y over an ecological context factor Z.

It takes some time and several examples, though, to understand this unusual conception. Imagine you were a teacher interested in the correlation between motivation (X) and achievement (Y). To assess the correlation, you need a sample of observation of x_i, y_i pairs, referring to motivation and achievement of different students *i*. Often, however, joint observations of both x_i and y_i in the same students are either not available, or separated in time, when achievement feedback (y_i) occurs with a long delay after the observation of motivation (x_i), rendering it hard to link the remote memory traces. Nevertheless, one often has baserate information that, say, most students in the class are high in motivation and also that most students are high in achievement, either in direct comparison to another observed class with mostly low motivation and low achievement, or in comparison to generalized knowledge about normal levels of motivation and achievement.

In any case, from the co-occurrence of elevated univariate baserates for X and Y, a positive bivariate correlation is inferred and used for predictions about individual students, even though the actual correlation across all x_i, y_i, pairs (had they been assessed) may in fact be zero or even negative. From this sketch of a PC illusion it can be seen that baserates are mistaken for a contingency or, with reference to a 2 x 2 table, the cognitive process relies on the marginal frequencies (baserates) rather than the joint frequencies in the cells of the table.

In the present article, we delineate the theoretical underpinnings of PCs and we review the pertinent empirical research. In the first section, we draw an analogy between PC inferences and the notion of ecological correlation, a phenomenon that has long intrigued the social sciences and that helps to understand PCs as correlations existing at a higher level of aggregation. Then, in a second major section, we review empirical evidence on PCs, emphasizing the generality of the phenomenon and addressing its paradigmatic constraints and potential moderators. A third section is devoted to a discussion of the adaptive value of PCs as a key heuristic to inductive reasoning in a complex and probabilistic world. In the final Conclusions section, we point out limiting conditions and we set PCs apart from related theoretical concepts and research paradigms.

Introduction to Pseudocontingencies

In a classical *Sociological Review* article, Robinson (1950) explained the phenomenon of ecological correlations with reference to race and education. When computed at the level of individuals, the correlation between race (White vs. Black) and illiteracy (literate vs. illiterate) within the US was negligible, in the range of .20. When computed at the level of ecologies (i.e., based on the *proportion* of Black people and the *proportion* of illiterates in different districts), in contrast, the correlation was substantial, exceeding values of .70 (ecological correlation across 48 US-States) or even .90 (ecological correlation across nine larger geographical divisions of the USA). Thus, the correlation of the same two variables – race and education – changed dramatically with aggregation level.

The correlation between the same two variables can vary strongly with the level of aggregation, because different causal factors exert their influence at different levels (Hammond, 1973). At a high aggregation level, the proportion of Black people may be large in districts with low housing prices, due to differences in income, and the same districts may produce a large proportion of illiterates, due to insufficient funding for public schools. In more prosperous districts, the proportions of Blacks and illiterates may be low. Thus, the district-level baserates reflect causal influences of education and income, which are independent of residents' genetic race (Fairchild, 1991), which might account for a correlation at the individual level. In fact, the elevated illiteracy rate in poor districts may be even higher among White people, meaning that the *partial* correlation between race and illiteracy, controlling for the impact of districts, may have an opposite sign.

There is no logical or rational basis to assume that individual-level correlations *per se* are more genuine or more correct than ecological ones. Assuming that either individual or aggregate correlations are normatively correct would be as meaningless as pretending that either partial correlations or zero-order correlations have general priority. Researchers examining genetic factors may be interested in the individual correlation between race and illiteracy (with differences between districts partialled out), whereas sociologists may be interested in the ecological correlation of race and illiteracy base-rates.

One might discard such ecological biases as rare and far-fetched, but discrepancies between the correlations observed in the same data at different levels of aggregation are actually quite common. On one hand, many research designs invoke multi-level data sets, with the same dependent measure assessed repeatedly over an extended series of trials, such as time series analyses, repeated measures designs, or designs involving hierarchically nested groups. On the other hand, many psychological phenomena constitute multi-level problems, such as the big-fish-little-pond effect in education (Marsh & Hau, 2003), the solution of hidden profile tasks in group decision making (Mojzisch & Schulz-Hardt, 2006; Stasser & Titus, 1985), or Simpson's paradox (Simpson, 1951). Social dilemmas (Chater, Vlaev, & Grinberg, 2007) provide a particularly prominent case. At the level of individual trials, defection is related to higher payoff than cooperation. On aggregate, however, over many trials, defection leads to lower payoff than cooperation.

Components of Covariance

To be precise in talking about correlations, we shall distinguish between an aggregate level (with K ecologies E_k, such as districts, groups or categories; $k = 1, 2, \ldots, K$) and an individual level (with n_k individuals or specific units, $i = 1, 2, \ldots, n_k$, within each ecology E_k). Given this two-level structure[1], three ways of computing a correlation between two variables X and Y can be distinguished:

(a) the *total* correlation $r(x_i, y_i)$ between individual X_i and Y_i scores, pooling across all $n_1 + \ldots n_k + \ldots + n_K$ data pairs from all K ecologies;

(b) the *ecological* correlation $r({}_{mean}X_k, {}_{mean}Y_k)$ at aggregate level of all K pairs of mean ${}_{mean}x_k$ and ${}_{mean}y_k$ values, across the K ecologies;

(c) and the *partial* correlation $r(x_i, y_i|E)$ between individual x_i and y_i controlling for the impact of ecologies, which means to compute the individual correlation within ecologies, and then to pool these within-category correlations across all K ecologies.

The three components of covariance resemble the familiar analysis-of-variance distinction between total variance, between-group variance, and within-group variance, respectively. Just as the variance components in the ANOVA, the component correlations can differ markedly. Crudely speaking, the total correlation can be considered a compromise of the ecological correlation across ecologies and the partial correlation within ecologies. The nature of this compromise depends on the overall variance accounted for by the two components. The compromise will be dominated by the ecological correlation when the differences between ecologies are large relative to the differences within ecologies (Figure 1a). It will approximate zero, if differences between ecologies equal differences within ecologies (Figure 1b). When both component correlations point in the same direction, they may still contribute different variance proportions to the total correlation (Figure 1c and d).

To be sure, there are constraints imposed on the relationship between total, ecological, and partial correlations (Nickerson, 1995; Pedhazur, 1982). Given a perfect positive correlation in both components, the total correlation also has to be positive. Over a wide range of naturally occurring correlations, however, the latitude for divergence is considerable.

From Ecological Bias to Pseudocontingency

PCs constitute a cognitive analog of ecological bias. Just as social scientists use aggregate data as a substitute for unavailable or unreliable individual data, so may laypersons derive estimates of individual level correlations from aggregate data whenever individuating data are unavailable in the environment, or in memory. The following PC effects obtained in a simulated classroom paradigm illustrate the phenomenon (Fiedler, Freytag & Unkelbach, 2007; Fiedler, Walther, Freytag, & Plessner, 2002).

Participants played the role of a teacher gathering observations about 16 students of a school class represented on a computer screen. In one experiment (Fiedler et al., 2007, Exp. 2), they had to teach a civics lesson in which all 16 students provided pro and anti statements about "Freedom of Science" (X) and "Benefits of Humanism" (Y), respectively, in two successive discussion rounds. The total correlation between the attitude domains was carefully controlled to be zero, $r(x_i, y_i) = 0$. For each attitude position in one domain, there was exactly the same number of pro statements and anti statements for the other domain (Table 1). However, the 16 students comprised two subgroups allegedly coming from different former teachers (or ecologies), marked by different colors. The ecological correlation between the average attitudes per subgroup was manipulated to be either positive or negative. In the positive PC condition (see the left panel of Table 1), the baserate of pro statements was high for both domains in one subgroup and low for both domains in the other, yielding a perfect positive ecological correlation of $r({}_{mean}X_k, {}_{mean}Y_k) = 1$.[2] In the negative PC condition, there was a perfect negative ecological correlation of $r({}_{mean}X_k, {}_{mean}Y_k) = -1$, because the subgroup baserate of pro statements was always high for one domain and low for the other (see the right panel of Table 1). Opposite to the ecological correlation, the partial correlation $r(x_i, y_i|E)$, that is, the correlation within subgroups, was –.58 in the positive PC condition and +.58 in the negative PC condition.

When the lesson was over, teachers were asked to rate the attitudes of each of the 16 students, prompted one at a time, in the two attitude domains. The school context made it perfectly clear that judgments of the individual students' attitudes were called for rather than judgments of average attitudes per subgroup. Although the total correlation over all individuals was clearly zero in either condition, the correlations between the 16 pairs of judgments were positive in the positive PC condition, but negative in the negative PC condition. When students belonged to subgroups with positively correlated average attitudes (at the aggregate level), their attitudes (at the individual level) appeared to be positively correlated, and vice versa, for students belonging to subgroups with negatively correlated average attitudes. Put differently, correlation inferences at the individual level were biased towards the correlation that held at the aggregate level.

The confusion of ecological and total correlation may be even stronger when there are more than two categories, or ecologies. Engage in the following thought experiment and imagine you were a teacher who, at the beginning of a school year, gets acquainted with four new classes (see Figure 2). You enter the first classroom, and there are 75% boys and 75% good students. In the second class, the proportion of boys is 25%, and the proportion of good students is 25% only. In the third classroom, again 75% are boys and 75% perform well. And in the fourth class, 25% are boys and 25% perform well again. Across classes, then, the ecological correlation between gender (X) and performance (Y) is perfect, $r({}_{\text{mean}}X_k, {}_{\text{mean}}Y_k) = +1$. Within classes, however, the proportion of good boys is lower than that of good girls, so the partial correlation is negative, $r(x_i, y_i|E) = -.33$. The total individual correlation between gender and performance, pooling across ecologies, is again $r(x_i, y_i) = 0$. Teachers exposed to such information may develop a rather strong belief that boys generally outperform girls (cf. Fiedler et al., 2007). In drawing this conclusion, however, teachers rely on the ecological correlation, although the task calls for the total individual level correlation.

Setting PCs Apart From Ecological Bias

Having introduced PCs as a cognitive analog of an ecological correlation, we hasten to add that our initial definition of PCs is broader and less restrictive than that of ecological correlations. For an ecological correlation to be observed, there have to be at least two ecologies. The cognitive PC illusion, however, generalizes to the case of a single ecology, for in the absence of an explicit contrast class a teacher can still have implicit background knowledge about a typical class. Enhanced baserates of boys and high performance in a single observed class, relative to such a standard of comparison, logically imply an ecological covariation of gender and performance baserates (across the observed class and the standard class). Two jointly skewed baserate distributions (or marginal distributions of a contingency table; cf. Figure 3a) are thus sufficient to induce a PC effect. If both distributions are skewed in the same direction, the resulting PC is positive, as is the inferred ecological correlation. If they are skewed in opposite directions, the resulting PC is negative. No distinct inference process has to be postulated to account for PCs arising in single-ecology settings. Although PCs may become stronger when explicit stimulus data on more than one ecology render the contrasting baserates salient, the same basic principle can account for PCs arising in tasks involving single or multiple ecologies.

Setting Pseudocontingencies Apart From Ordinary Contingencies

The special case of PCs resulting from skewed baserates in a single ecology, which does not constitute a full-fledged ecological correlation, is also important to set PCs apart from contingencies proper. An ordinary contingency implies an interaction: The conditional rates of high and low *Y* values given *X* should change with *X*-levels. If there are but two main effects, say, a tendency toward high *X* and high *Y* values, but no interaction (i.e., the rate of high *Y* values is constantly high for all values of *X*), the contingency is zero. In contrast, PCs reflect an erroneous inference that mistakes two main effects for an interaction, as if the observation of high baserates of both *X* and *Y* implied an increase in *Y* from one level of *X* to the other. Figure 3 shows that this is not the case; the co-existence of two skewed base-rates (e.g., mostly boys and mostly high performance) allows for a zero (Figure 3b), a positive (Figure 3c), or a negative (Figure 3d) correlation.[3]

An even more striking difference between PCs and ordinary contingencies originates in their fundamentally different computational operations. For a contingency between *X* and *Y* to be assessed, there have to be joint observations of bivariate x_i, y_i, pairs. In the example, the teacher's observations have to reveal both the gender and the performance of a student. In reality, though, information about the co-occurrence of x_i and y_i is often not available. A teacher may observe a high baserate of male students in a class on one occasion (e.g., while screening a list of students' names) and a high baserate of good performance on another occasion (e.g., while correcting homework, performance should be more salient than gender). In such a learning environment the contingency remains undefined (cf. Figure 3a), but the baserates or marginal distributions still allow for PC-type inferences. The cognitive inference process is thus crucially different from genuine correlation assessment.

At the operational level, this crucial difference can be manipulated through different presentation modes. In an experiment using a *simultaneous* presentation format, which allows for contingency assessment proper, bivariate information about two stimulus aspects, x_i and y_i, is presented simultaneously. For example, information about gender and performance would be presented together, jointly referring to the same student. By contrast, in a *successive* presentation format, information about gender (X) would be presented in a first run, and information about performance (Y) would be presented subsequently in a second run, so that the joint reference of values x_i and y_i to the same student i is not easily apparent.

One might wonder whether under such conditions participants refuse to judge contingencies, but research on PC illusions shows that they readily infer contingencies when no joint observations are available. PC effects obtained under such task conditions are clearly distinct from the large number of studies on contingency assessment and illusory correlations (Allan, 1993; Alloy & Tabachnik, 1984; Crocker, 1981; DeHouwer & Beckers, 2002; Fiedler, 2000; Hamilton, 1981; White, 1995), because the successive presentation format rules out the computation of contingencies proper.

Having pointed out the distinct nature of PCs – in situations in which contingencies are undefined – we hasten to add that PC effects may, to an unknown degree, intrude into the simultaneous presentation paradigm of contingency assessment. That is, even when the joint frequencies of all x_i, y_i combinations (i.e., the cell entries) are given, the marginals or baserates are available too. It is uncertain to what extent the cognitive process on such simultaneous tasks is driven by the joint frequencies (contingency) or by the marginal frequencies (pseudocontingency). Prior contingency research has presupposed that subjects do what statisticians in the Spearman-Brown tradition expect them to do, but our research suggests that they may silently treat such tasks like PC inference tasks, utilizing baserates in addition to, or instead of, cell frequencies. This can be demonstrated in experiments in which the stimulus series is constructed to yield a PC that diverges, in size and/or in sign, from the actual contingency. Such an example appears in Figure 3d. Although the univariate distributions are skewed in the same direction, supporting a positive PC between *X* and *Y*, the cell entries reveal a negative contingency (i.e., the proportion of high *Y* values is lower for high than for low *X* values in Figure 3d). When PC effects in such situations override divergent contingencies, this provides especially strong evidence for their distinct nature.

The Illustrative Case of Illusory Correlations

To illustrate the competition of PCs with genuine contingency information, consider the host of findings on the phenomenon of *illusory correlations* (for an overview, see Fiedler, 2000). In one prominent paradigm, a correlation is perceived in a stimulus design with two binary variables that have skewed marginal distributions but are stochastically independent in the stimulus set. Usually the two variables are group membership, with Group A being more frequent than Group B, and desirability of behavior, with desirable behaviors being more frequent than undesirable behaviors (Hamilton & Gifford, 1976). Despite the fact that the ratio of desirable to undesirable behaviors is identical for the majority and the minority, the majority is judged more positively than the minority. This illusory correlation effect was originally explained as reflecting the enhanced availability in memory of events that combine the infrequent group membership (i.e., Group B) with the infrequent class of behaviors (e.g., Hamilton & Gifford, 1976; Stroessner, Hamilton, & Mackie, 1992).

Traditional accounts of illusory correlations tacitly assume that the memory representation underlying biased group judgments consists of individual behavioral incidents, along with their group origin. Contrary to this assumption, however, memory analyses of discrimination performance concerning the group origin of desirable and undesirable behaviors in the illusory correlation paradigm not only failed to show a memory advantage of paired-infrequent events (Fiedler, Russer, & Gramm, 1993; Klauer & Meiser, 2000), but also revealed that source memory for the group origin often approaches zero for all combinations of group membership and desirability (Meiser, 2003; Meiser & Hewstone, 2001, 2006). Hence, the memory representation does not preserve sufficient information about the joint occurrences of Group A and Group B with desirable and undesirable behaviors to allow for a (biased) memory-based assessment of the correlation.

As an alternative to the traditional account on the basis of memory for individual co-occurrences, the PC framework can easily accommodate the illusory correlation effect as an inference on the basis of skewed marginal distributions, which can be extracted quite accurately from a 2x2 stimulus design (Reips & Waldmann, 2008). According to this interpretation, the majority group is aligned to the frequent class of desirable behavior, and the minority group is aligned to the infrequent class of undesirable behavior. The PC interpretation of illusory correlations gains support from earlier experiments using only baserate information, rather than information about the joint occurrence of group membership and behavior (McGarty, Haslam, Turner, & Oakes, 1993): Simply telling participants that there are twice as many statements about Group A as Group B, together with the presentation of more desirable than undesirable behaviors without group labels, was sufficient to elicit a pronounced illusion in favor of Group A. Together with the aforementioned memory results, these findings demonstrate that illusory correlations may not reflect information about the interior of the contingency table, or on enhanced memory of information from one cell of the table. Instead, illusory correlations may reflect inferences from baserates.

A feedback learning variant of the same task highlights this dominance of baserates over joint-occurrence rates. In each trial of a study by Eder, Fiedler, and Hamm-Eder (2008), participants who had been told that Group A was larger than B were asked to predict the valence of a behavior associated with a given group. Feedback was then provided about the valence and content of the actual behavior. Over many trials, the feedback revealed that most behaviors were positive. Participants easily learned to predict mostly positive behaviors, but this bias was stronger for Group A than Group B. These baserate-driven PC predictions overrode the genuine contingency information entailed in the corrective feedback, which continuously reminded participants that Group A behaviors were less positive than predicted.

Understanding PCs From an Environmental-Learning Perspective

Why do judges and decision makers, or organisms more generally, apply the cumbersome PC inference scheme and often even prefer it over the contingency that statisticians have taught us to be the normative model? The apparent category mistake underlying PC illusions notwithstanding, a moment of reflection on its origins reveals a number of plausible reasons why PCs may be more useful and more functional a heuristic than contingencies proper. Although statistics courses generally start from a matrix involving N subjects and M variables, nature would rarely provide us with such fully connected multivariate data. Reconsider the case of a teacher who is to detect or assess antecedent conditions of students' performance. Among these antecedent conditions are such diverse variables as gender, intelligence, achievement motivation, interest in a discipline, socio-economical status, TV consumption, former teachers, teaching methods, school system variables, and many others. During the long time the teacher is observing the students, she does not know which specific contingency hypothesis she may have to test in the future. Will performance be related to TV, or to motivation, or to status? Teachers (or learning organisms) must be prepared for all kinds of contingency problems that may arise in the future.

If the teacher were to frame this demanding task as a genuine contingency problem, he or she would have to assess a complex frequency table, covering the combinations of all levels of all variables potentially related to student performance. The number of cells of such a table can easily exceed a manageable range. Assuming only binary variables, the function 2^N, the total number of cells of the resulting frequency table, already exceeds one-hundred for only $N = 7$ variables. If more than two variable levels have to be distinguished, the resulting huge contingency table would be impossible to keep in memory for the sake of judgment formation. Even with the help of external storage devices, the cells would be replete with missing data, because nature does not provide us with complete multivariate vectors that specify the value of every variable for every observation. Thus, the teachers' environment does not provide complete information about all students' values on all variables.

Rather, the learning environment in which organisms must observe the correlates of any variable (e.g., performance) typically provides us with information about one, or very few, variables at a time, and observations about the same object are separated by countless observations about other objects. On one occasion, the teacher learns about performance and motivation, then about performance and gender, and on still another, detached occasion about performance and TV consumption. In fact, it may often be the case that organisms encode only one variable at a time. That is, the teacher may assess students' performance in physics in a physics lesson, the same students' performance in chemistry in another lesson, she learns about their sociability on an excursion, and about their family background during a parents meeting. To relate performance to such variables, the teacher has to coordinate in memory all these separate episodes. This extremely taxing and demanding memory task would easily exceed the memory capacity for attribute-student relations.

In fact, nature is not only reluctant to provide joint information about all multivariate measures at the same time. Such information is often not available at all, or completely inaccessible at a given point in time. Learning about the family background and the living conditions of students may have taken place long ago. Feedback about the outcomes and consequences of actions and decisions often appears, if at all, only with a long delay. When new variables are finally discovered, all the other variables would have to be broken up by the new variable, if the ideal of the full multivariate array had to be maintained. Thus, for many different reasons, a full-fledged contingency table does not describe the learning environment of organisms in an uncertain world. The actual learning environment, rather, would appear to be more compatible with the simple baserate knowledge that underlies the PC algorithm than with the ideal of statistics textbooks that take multivariate data points for granted.

Review of Empirical PC Research

This section provides a review of empirical studies that explicitly refer to the term "pseudocontingency". Later on, we will also mention implicit PC studies that were conducted with a different rationale in mind but that can be interpreted, post hoc, as instances of PCs. Our empirical review has several aims: to demonstrate the pervasiveness of PC illusions across a wide range of task contexts and materials; to describe various methods providing convergent evidence for PCs; and to identify boundary conditions that reveal the nature of the underlying cognitive process.

Alignment of Baserates

Theoretically, PC effects should arise when in a distinct ecology (i.e., category or group) the marginal distributions or baserates of two attributes are skewed, either in the same direction (yielding positive PCs) or in opposite directions (yielding negative PCs). Psychologically, this basic prediction should only hold when two boundary conditions are met, namely, (a) that participants notice the skewed baserates; and (b) that PC effects are not overshadowed by other processes. The first boundary condition (a) implies that PCs should correlate with participants' awareness of the skewed baserates. The second condition (b) means that pure PC effects should be most likely when both prior expectancies and genuine contingency information about the joint occurrences of X and Y are ruled out although, as already mentioned, PC may even survive counteracting contingencies.

PCs ruling out both prior expectancies and genuine contingencies. An experiment by Fiedler and Freytag (2004, Exp. 1) represents a paradigm developed to meet both boundary conditions. Participants were forced to encode two skewed stimulus attributes (a), and a successive presentation mode was used to eliminate genuine contingency information (b). Participants were first presented with information about the diets, vegetarian or prebiotic, assigned to patients in two different wards (ecologies) of a hospital. In one ward, A, a majority of 36 out of 48 patients received a prebiotic diet and only 12 patients were treated with a vegetarian diet. In the other ward, B, the diet baserates were skewed in the opposite direction (36 vegetarian and 12 prebiotic). Each stimulus display presented a patient's 4-digit ID number, the ward to which he had been assigned and the diet applied during the past week. To warrant accurate encoding of the diet information, participants had to indicate, for each patient, whether his diet plan was dominated by prebiotic or vegetarian food. In a second run, participants learned about the number of symptoms observed in each patient. Each stimulus display again consisted of the patient's ID number and the ward membership plus a seven-row table showing the number of symptoms observed in the patient over the seven days of the week. Participants again had to type in their estimate of the average number of daily symptoms. This variable was also skewed within both ecologies (see Table 2a). In ward A, the numbers of patients with many and few symptoms were 36 and 12, respectively; in ward B, the baserates were reversed, 12 patients having many and 36 having few symptoms. Thus, the alignment of baserates in both wards – mostly prebiotic diet and many symptoms in ward A and mostly vegetarian and few symptoms in B – should produce a PC linking prebiotic diet to high number of symptoms. If we call this relation positive, then a negative PC (linking prebiotic diet to low prevalence of symptoms) should emerge if the symptom baserates of the two wards were reversed (i.e., few symptoms and mostly prebiotic diet in ward A, many symptoms and mostly vegetarian in B). In a control condition, baserate distributions for symptoms and diets were equal (i.e., 24 patients each for every diet and symptom level).

The actual correlation between dieting (X) and symptoms (Y) was unknown, because the successive presentation format kept participants from coordinating the X and Y values pertaining to the same patient ID number. This did not prevent them from using one variable when making predictions of the other, though. When, in a subsequent prediction task, they were either given a patient's diet and asked to predict the same patient's symptom frequency, or when they were given a patient's number of symptoms and asked to predict the same patient's diet, the correlation they produced between predictor and prediction was significantly positive in the positive PC condition, negative in the negative PC condition, and close to zero in the control condition. Table 3 shows the average estimates of the predicted variable as a function of different levels given on the predictor variable. Apparently, judges readily used one variable to predict the other. The absence of prior expectancies is evident in the zero effect for the control condition. Depending on whether the baserates of prebiotic diet and the baserates of high symptom frequencies pointed in the same direction, or in opposite directions, the predictions reflected either clearly positive or clearly negative relationships between the two variables.[4]

Another subsidiary finding is telling about the cognitive process underlying PCs. Given the positive PC condition, with 75% prebiotics and 75% high-symptom patients, but only 25% vegetarians and 25% low-symptom patients in a ward (cf. Table 2a), predictions for the rare events reflected the same positive correlation as predictions for the frequent events. That is, few symptoms were as readily associated with vegetarian diet as many symptoms were associated with prebiotic diet (cf. Table 3). This remarkable finding means that PCs do not just reflect the number of associative-learning trials or the association strength of specific stimulus pairings. They rather reflect rule-based inferences from the alignment of baserates.

PC effects overriding divergent contingencies. In other experiments reported by Fiedler and Freytag (2004, Exp. 2 and Exp. 3), participants were informed about the level of various respondents belonging to one of two patient groups, A and B, on two fictitious personality tests, X and Y, denoted abstractly to rule out prior expectancies. In one group, A, there were mainly high X scores and mainly high Y scores; in the other group, B, there were mainly low scores on both tests, X and Y.

Successive versus simultaneous presentation format was varied. Given a successive presentation mode, the presentation of all X scores was complete before the Y scores of the same respondents were given; so there was little chance to coordinate X and Y scores belonging to the same persons in memory. However, in the simultaneous presentation condition, when the same respondent's X and Y score were presented adjacently, the task setting clearly encouraged subjects to assess the contingency proper. The (partial) correlation within both groups was exactly zero, $r(x_i,y_i|E) = 0$, as the ratio of high and low values on one test was equally likely for high and low values on the other test (i.e., 16:8 = 8:4). The skew of the resulting baserates was always 24 versus 12 (cf. Table 2b). The total correlation, when the 2x2 tables for both groups were pooled, was slightly positive, $r(x_i,y_i) = +.20$. The test persons' group membership was indicated by different background colors of the screen.

Again, when participants were subsequently asked to predict new patients' Y test scores from their X test scores, or vice versa, the predictions regularly exhibited a positive relation. For both presentation modes, the PC effect was significantly above zero (the partial correlation) and clearly higher than +.20 (the total correlation). PCs were slightly more pronounced for the successive than for the simultaneous presentation mode. This makes sense because successive presentation not only eliminates competing contingency information but should also facilitate the encoding of X and Y baserates. PCs in the successive-presentation condition indicate that pseudocontingencies are readily derived from baserates when genuine correlations are undefined. PCs after simultaneous presentation highlight that baserate-driven PC effects override the actually present correlation when joint x_i,y_i observations are available. This means that the cognitive process may follow the marginal distributions more than the cell entries of the contingency table.

Even more impressive PCs overriding contingency information were obtained (Fiedler & Freytag, 2004, Exp. 4) with three score levels (low, medium, and high) on test X and Y. The 3x3 contingency scheme (Table 2c) yielded a zero total correlation, $r(x_i,y_i) = 0$, and a partial correlation within groups that was clearly negative, $r(x_i,y_i\,|E) = -.43$. The ecological correlation to be sure, given two groups, was perfectly positive, $r({}_{mean}X_k, {}_{mean}Y_k) = 1$ (see Footnote 2). Predictions of individual persons' scores on one test from a given value on the other test reflected strongly positive correlations (> +.50 for three quarters of all participants).

This positive PC effect was equally strong for the group with mostly high *X* and mostly high *Y* values as for the group with mostly low values on both tests. Again, on the prediction test, judges not only associated frequently observed combinations (e.g., high *X* and high *Y* scores in Group A). Rather, when given a low *X* score of a person from Group A, which had never (!) been paired with a low *Y* score in the stimulus series (see Table 2c), they nevertheless predicted a low *Y* with the same regularity as they predicted a high *Y* score from a high *X* score (and vice versa). This nicely corroborates that PC effects do not reflect the associative strength of individually learned stimulus pairs but the alignment of dimensional baserates, extending to combinations that have never been observed.

Common causes and common effects. PCs also generalize across different causal models that might explain, or rationalize, inferences from ecologies to individual scores. Fiedler and Freytag (2004, Exp. 2 and Exp. 3) told their participants that the two groups A and B referred to persons that did or did not undergo psychotherapy. In the common-cause condition, they were told that psychotherapy may have a causal influence on both tests, *X* and *Y*. Such a causal model, which suggests that *X* and *Y* reflect in part the same causal influence, would actually justify the inference that individual *X* and *Y* scores must be correlated across groups (Waldmann & Hagmayer, 2001). In contrast, in the common-effect condition, they were told that *X* and *Y* scores jointly determine whether patients have to undergo psychotherapy. Such a causal model provides no reason why *X* and *Y* scores should be correlated. PC effects were obtained in both conditions, although the illusion was slightly stronger in the common-cause than in the common-effects condition.

Singular and multiple ecologies. The studies reviewed so far used two contrasting categories or ecologies with opposite baserate distributions. This task setting might be particularly PC-prone, because it maximizes the impact of the ecological correlation, which is always perfect (i.e., $r({}_{\text{mean}}X_{\text{k}}, {}_{\text{mean}}Y_{\text{k}}) = 1$ or $r({}_{\text{mean}}X_{\text{k}}, {}_{\text{mean}}Y_{\text{k}}) = -1$) for two ecologies (see Footnote 2), whereas individuating correlations are often complex and noisy. Indeed, when comparing two-category PCs to single-category PCs within the same experimental task (Fiedler & Freytag, 2004; using the stimulus distribution of Table 2c), the resulting correlations between given and predicted test values were on average .60 and .71 (for the high and low scoring groups, respectively), as compared with only .19 in a condition that only received information about one group.

Moving from two-ecology problems to multiple-ecology problems, though, ecological correlations may be less clear-cut and more difficult to discern, but PC effects may arise for different reasons. The complexity and information overload of such problems may preclude the assessment of individuating data from multiple categories and force judges instead to form a more economical code at categorical level. To illustrate the impressive task of assessing information from as many as 16 categories simultaneously, consider the following variant of PCs in a simulated classroom setting (Fiedler et al., 2007, Experiment 4). In this study, a class of 16 students (8 boys and 8 girls) was represented graphically on the computer screen. Participants played the role of a teacher who, over a long list of trials, asked knowledge questions, observed which students raised their hands, selected one student, and assessed whether his or her answer was correct or incorrect. Each student had a motivation parameter (i.e., the probability of raising the hand) and an ability parameter (i.e., the probability of giving a correct answer). The participant's task was to assess and evaluate all 16 students' performance after each lesson, providing percentage estimates of their motivation (% raising hand) as well as their ability (% correct answers).

The actual correlation between motivation and ability was always zero. At the level of students, the motivation parameters varied from .2 (4 students) to .5 (8 students) to .8 (4 students), but all 16 students' ability parameter was held constant at .7 (in one class) or at .3 (in another class), thus precluding any correlation. At the lower aggregation level of responses to individual questions, the computer's random generator also precluded any correlation; each student's probability of a correct response on a particular trial was the same regardless of whether the student had raised his or her hand or not. Thus, only the ability and motivation baserates varied between students, whereas all correlations were necessarily zero, within students, and pooling all responses from all students.

However, crucially, within each student considered as a separate ecology or category, the conditions for a PC effect were met. In a class with a constantly high ability baserate, a positive PC between raising hands and correct responding should be expected for students with a high motivation base-rate, but a negative PC for students with a low motivation baserate. Conversely, in a low-ability class, a negative PC can be expected for high motivation students and a positive PC for low-motivation students. In accordance with this prediction, in high-ability classes, but not in low-ability classes, teachers' judgments of motivation (hand raising) and ability (correct responses) were correlated. Presumably, the complexity of the task to assess ability and motivation in as many as 16 different categories rendered the memorization of elementary behaviors impossible and forced judges to form an economical code for motivation and ability at the super-ordinate level of student categories.

Evidence for a Distinct Cognitive Process Underlying PC Effects

The purpose of the following section is to provide evidence for an aggregate-level cognitive process supposed to underlie the PC illusion. Most pertinent evidence comes from experiments that include multinomial-modeling methods to point out judges' sensitivity to differential baserates and their relative insensitivity to individuating raw information.

Evidence for aggregate-level representations. Several experiments focused on PCs in the area of stereotype formation and analyzed the kinds of memory representation that contribute to biased judgments about social groups (Meiser, 2003, 2006; Meiser & Hewstone, 2004, 2006). In an extension of the already introduced paradigm, participants were presented with a series of desirable and undesirable behaviors performed by members of two social groups, labeled Group A and Group B, in two different towns, labeled Town X and Town Y. Each behavioral description was displayed together with group membership and town context.

Thereby, the stimulus set provided information about the baserates of towns, groups, and desirability. Group A formed the majority and desirable behaviors were more frequent within Town X, whereas Group B formed the majority and undesirable behaviors were more frequent within Town Y (see Figure 4a). The correlation between group membership and desirability, which is of main interest for group evaluations, differed between the levels of aggregation. While the partial correlation within each town was zero (Meiser & Hewstone, 2006) or even negative (Meiser, 2003; Meiser & Hewstone, 2004), the correlation was positive when aggregating across towns. Judgments of the two groups in the two towns were assessed by means of trait ratings, frequency estimates of undesirable behaviors, and assignments of desirable and undesirable behaviors to the two groups within the two towns in a source-memory recognition task. All measures consistently revealed a more positive evaluation of Group A than Group B, reflecting the inference of a positive correlation of groups (A vs. B) with desirability, in spite of the actual zero or even negative correlations within each town context. Multinomial analyses of recognition and source memory for the town and group origin of the behaviors provided converging evidence that the resulting bias in favor of Group A was based on memory for aggregate-level information rather than exemplar memory for individual events, as we shall review next.

Evidence against spurious correlations. Biased group judgments with a confounding context factor, like the town factor in the present scenario, have previously been explained in terms of Simpson's paradox, assuming that judges ignore the role of the context variable, towns, and merely attend to the covariation of groups and valence (Schaller, Asp, Rosell, & Heim, 1996; Schaller & O'Brien, 1992). According to this explanation, the observed stereotype in favor of Group A reflects the spurious total correlation between group membership and desirability, which emerges when people fail to take into account the role of the town variable. This interpretation, however, does not apply to the present findings. The results show (a) that people do not rely on the co-occurrence of group membership and valence and (b) that they do take the ecological context factor into account.

Pertinent evidence about the memory representation used for the group judgments comes from analyses of episodic and reconstructive memory by means of multinomial models (Batchelder & Riefer, 1990, 1999). Specifically, a multinomial model of multidimensional source memory (Meiser & Bröder, 2002) was applied to the responses in an assignment task, which required subjects to first make recognition decisions about desirable and undesirable behaviors and, if they recognized a behavior, to recall the town and the group with which it had been associated. This analysis affords measures of source memory for the group origin and the town context of recognized behaviors. Source memory for the group origin of individual behaviors was very poor and not significantly different from chance.

Thus, despite the fact that the simultaneous presentation format provided immediate information about the co-occurrence of group membership and behavior, the memory representation at the time of judgment usually did not retain the group origin of behaviors. This finding speaks against an interpretation in terms of a spurious total correlation, for which paired information about group membership and behavioral valence is a prerequisite.

Importantly, even when source memory for the group origin of desirable and undesirable behaviors was raised to a substantial level through specific encoding manipulations, this did not affect the strength of the stereotype bias (Meiser, 2003; Meiser & Hewstone, 2004). More specifically, the stereotype was not strengthened when deliberate manipulations enforced a memory representation that binds behavior and group information together, which is inconsistent with a spurious correlation between groups and desirability that ignores the town context variable. These findings strongly suggest a cognitive process that does not utilize the joint frequencies defining the contingency proper, but other aspects of information that can be extracted from the stimulus series, such as the varying baserates of desirability and of the two groups across the town ecologies. [5]

Another finding that speaks against an explanation in terms of a spurious total correlation arising from a neglect of the context factor is that participants consistently showed sensitivity to the correlations of the ecological variable (towns) with desirability and group membership in their judgments (Meiser, 2003; Meiser & Hewstone, 2004, 2006). Overall, group members in Town X were evaluated more favorably than group members in Town Y in trait ratings and frequency estimates, reflecting the actual difference in the two towns' desirability (see Figure 4b). Nevertheless, the estimated proportions of undesirable behaviors were higher for Group B than Group A within the two towns, reversing the actual decrease in the proportion of undesirable behaviors from Group B to Group A.

In a similar vein, reconstructive guessing biases in the source-memory decisions showed that participants noticed and utilized the existing correlations between towns and desirability and between towns and groups: As illustrated in Figure 4c, desirable behaviors were more likely to be assigned to Town X relative to undesirable behaviors, mirroring the true preponderance of desirable statements in Town X and of undesirable statements in Town Y. Moreover, desirable and undesirable behaviors in Town X were more likely assigned to Group A (A|X in Figure 4c) than were behaviors assigned to Town Y (A|Y), mirroring the true correlation between the groups and towns. Nevertheless, the rate of behaviors assigned to Group A was higher for desirable than undesirable behaviors, reversing the true proportions of Group A among desirable and undesirable statements. Last but not least, the noticeable overestimation of zero frequencies and the corresponding underestimation of true proportions of 100%, in Figures 4b and 4c corroborate the conclusion that the group judgments did not rely on the retrieval and integration of individual-exemplar information, but on aggregated representations of baserates in the trivariate correlation structure.

Together, the results in Figure 4 show that the formation of a biased stereotype coexists with the extraction of the true ecological associations of town context with both desirability and group membership. Further analyses revealed that awareness of the role of the ecological factor was positively related to the strength of the stereotypical bias across individual participants (Meiser & Hewstone, 2004, Experiments 1 and 2), resembling the relation between awareness of skewed baserates and the perception of positive and negative PCs obtained in other studies (Fiedler & Freytag, 2004). These results underline that the biased group stereotype does not reflect a spurious total correlation between group membership and desirability, resulting from a neglect of the context variable. Rather, the stereotype that favors Group A over Group B reflects a PC derived from varying proportions of Group A versus B and desirable versus undesirable behaviors across the two towns.

Isolating baserate differences from contingencies. The interpretation in terms of a PC was further corroborated in two recent studies (Meiser, 2006), using a zero total correlation between group membership and desirability pooled across the two towns, but a negative partial correlation (i.e., a higher proportion of desirable behaviors for Group B than Group A) within each town context (Figure 5, full-information condition). Given the uniform distribution in the aggregate table collapsed across towns, a perceived positive contingency that favoring Group A over Group B could not be due to a spurious correlation that is only sensitive to group-behavior pairings while neglecting the town context. The PC framework, in contrast, predicts that the alignment of skewed baserates – mostly Group A members and mostly desirable behaviors in Town X, but mostly Group B members and mostly undesirable behaviors in Town Y – should result in a higher degree of desirability associated with Group A than with Group B. The empirical results in fact confirmed this prediction, despite the zero total correlation and the negative partial correlations within each context.

According to the PC rationale, a bias in favor of Group A should not depend on whether the focal correlation of group membership and desirability is actually presented. The bias should also be observed if incomplete trivariate stimulus information is presented, providing only information about the differential baserates of Group A versus Group B and of desirable versus undesirable behaviors between Towns X and Y. Similar to the successive presentation format employed in Fiedler and Freytag (2004), the study by Meiser (2006) was therefore based on the incomplete-information information condition depicted in Figure 5, which fully precluded the assessment of a total or partial correlation between group membership and desirability, because no information about their co-occurrence was given.[6] Nevertheless, a positive PC was observed, reflecting more favorable judgments of Group A than Group B.

More Paradigms for PC Effects

All studies that have been reviewed so far were conducted with the explicit goal in mind to demonstrate and illuminate PCs proper. In this section, we move from explicit PC research to implicit PC evidence from studies that were originally motivated by different theoretical goals. This evidence, too, can be easily assimilated to the PC framework.

PCs Between Continuous Variables. The experiments reviewed so far have examined PCs between qualitative and mostly dichotomous variables such as target affiliation (Group A vs. Group B), valence of behavior (positive vs. negative), performance (good vs. poor), or symptoms (many vs. few). Many psychological phenomena, however, are rooted in continuous variables, such as intelligence, temperament, or sociability. Nevertheless, given sufficiently pronounced trends towards the high or low end of such continuous scales, the same alignment process is possible as with dichotomous variables. A positive (negative) PC between two variables can be expected when two trends in an ecology point in the same (different) direction. As in the case of dichotomous variables, the co-occurrence of trends at aggregate level should guide the cognitive process.

For a single-ecology example, Lambert and Wyer (1990) found that a target person (dis)confirming a stereotype in one aspect was expected to (dis)confirm the stereotype in other aspects as well. For another example involving two ecologies or groups, consider the following study by Freytag (2003). Participants first learned about the sociability and the studiousness of the members of two types of students, and were later asked to infer new exemplars' value on one dimension from knowledge about their value on the other dimension. As the participants saw the individual students' values on both dimensions in a simultaneous presentation task, the actual intragroup correlation, which was always zero, could be compared with the correlation expressed in participants' inferences.

The stimulus series always contrasted a focal group of students, which exhibited intermediate levels on both attributes, to a highly sociable context group. The only experimental manipulation consisted in the context group's position on the studiousness dimension, which was either high or low. Thus, given the constantly high position of the context group on sociability, the ecological correlation between the attributes was positive in the former condition (i.e., studiousness high in the context group), but negative in the latter condition (i.e., studiousness low in the context group). To the extent that inter-attribute inferences reflect a PC, positively correlated inferences should be obtained when the context group scored higher on both dimensions, but negatively correlated inferences should be obtained when the context group scored higher on one dimension but lower on the other.

As expected, inter-attribute inferences were substantially correlated in the direction implied by the PC rationale. These findings obtained although participants accurately reproduced the means and standard deviations for both attributes within both target groups. It thus seems unlikely that the group membership of the individual students had been ignored or forgotten. Rather, the joint observation of accurate univariate estimates per group and context-dependent inter-attribute inferences lends further support to the idea that PCs originate in the effective encoding of aggregate-level, ecological trends.

PCs in operant conditioning tasks. Using an operant learning paradigm modeled after Goodie and Fantino (1996, 1999), Kutzner, Freytag, Vogel, and Fiedler (2008) studied predictions of binary outcomes given one of two audio signals. Participants received outcome feedback, gaining and loosing money for correct and incorrect predictions, respectively. One outcome occurred clearly more frequently than the other, and one signal occurred more frequently than the other, while the outcome rates were not contingent on the preceding signal. Participants developed an adaptive bias toward predicting the more frequent outcome, which increased their payoff. However, crucially, the bias toward the more common outcome was more pronounced after the more common signal, reflecting a clear-cut PC strategy. Apparently, then, PC effects may even generalize to the domain of operant conditioning.

PCs Based on Propositional Information. That PCs can be derived from generalized baserate tendencies, independently of a specific stimulus format, is evident from the fact that even findings from reasoning experiments that use a propositional stimulus format can be paraphrased as PC effects. The well-known atmospheric effect in syllogistic reasoning (Klauer, Musch & Naumer, 2000; Woodworth & Sells, 1935) bears an obvious resemblance to PC-type inferences. Given two verbally expressed baserates in the premises “Most X are Z” and “Most Y are Z”, participants tend to believe that the conclusion “Most X are Y” is correct. Let X and Y be two attributes in an ecology Z, the analogy between PCs and atmospheric effects is striking. The convergence between inductive inferences and syllogistic reasoning highlights the contention that PCs follow a reasoning process based on an aggregate memory code, distinct from correlation assessment based on specific joint observations.

In a lecture hall demonstration, students learned that the majority of people living in a fictitious country called Xanadu were hunters and that the majority was often suffering from stomach aches. When asked to estimate separately the proportion of hunters and non-hunters, respectively, suffering from stomach aches, most students expected this proportion to be larger for hunters (Fiedler & Freytag, 1999). Finally, the McGarty et al. (1993) findings cited earlier offer yet another example of PC inferences from propositional information. PCs in propositional-reasoning tasks highlight the fact that a subjective sense of validity surrounds baserate-driven inferences, independently of the statistical complexity of the task, much like the sense of validity underlying erroneous syllogistic conclusions.

Summary of research review. In summary, the evidence reported in this section entailed some notable insights about the cognitive process that underlies the formation of PC illusions. Multinomial-model analyses have shown that the memory representation of the trivariate distribution of two correlated attributes across an ecological context variable may often lack systematic traces of the co-occurrence of particular attribute values. While the parameters for source memory, which constitutes the basis for contingency assessment proper, hardly differ from chance, the parameters for biases reflecting the baserates of attributes in different ecologies are substantial and predictive of the strength of the PCs. Even when enhanced source memory for the joint occurrence of attribute pairs is enforced experimentally, the perceived relation between attributes continues to depend on the ecological baserates. PC inferences can occur even when the two focal attributes were never paired in the stimulus series. The demonstration of PCs for continuous variables varying on a quantitative scale corroborates the contention that the cognitive process does not rely on the frequencies of particular event combinations. Rather, the crucial condition for a PC illusion seems to be the presence and encodability of an ecological correlation between attribute baserates, regardless of whether it is observed across two or more stimulus ecologies or inferred from a comparison of a single ecology with a prototype or standard ecology.

Adaptive Functions of PC Inferences:
Tracing the Illusion Back to a Rational Inference Scheme

So far, we have reviewed studies that stress the illusory character of PCs. Yet, in the initial discussion of the origins of the PC illusion, we have already provided arguments to justify the illusion, pointing out that there is hardly any alternative to PC inferences in a world that rarely provides organisms with multivariate contingency data. In this section, we will even go one step further and argue that PCs are not only a regrettable by-product of a nasty probabilistic world (Chater & Oaksford, 2008). Rather, we will reframe PCs as an adaptive inference algorithm, which must not be considered irrational.

A Learning Algorithm Leading to PC Inferences

Let us take the perspective of a rational-minded organism searching for an inductive-learning algorithm that renders the empirical world predictable and controllable. We will see that such a rationally motivated organism will often end up making inferences of the PC-type.

Defining adaptive learning. Assume that the goal of adaptive learning is to maximize correct predictions on a criterion variable, Y. For an animal, Y could be the occurrence of positive (food) or negative (predator) stimuli; for a teacher (i.e., a human judge), the criterion might be correct versus incorrect student answers. For simplicity, we assume a dichotomous criterion, Y_{frequent} versus $Y_{\text{infrequent}}$, with unequal occurrence baserates, $p(Y_{\text{frequent}}) > p(Y_{\text{infrequent}})$. Adaptive learning means to extract and utilize such information for predictions.

Functionality of response biases. In the absence of other information, maximizing the rate of correct decisions means to invariantly predict the more prevalent Y value. To detect changes in the environment, though, organisms often engage in probability matching rather than maximizing, predicting the more likely outcome Y_{frequent} at the same rate as it occurs, $p(Y_{\text{frequent}})$. In any case, a response bias toward the more likely criterion outcome would be functional (cf. Kareev, Fiedler & Avrahami, in press; Shanks, Tunney & McCarthy, 2002).

Using ecological and individuating predictors. Universal baserates, however, are virtually never known, nor would they be very useful for predictions in particular situations. Thus, the overall prevalence of a predator across time and space is of little value for the animal. What is needed, rather, is the conditional occurrence rate of the criterion event given specific predictor conditions. We can distinguish more general predictor conditions E_k and more specific conditions i. Let us refer to E_k as ecologies and to i as individuals; for instance, E_k could be the animal's territory and i could refer to specific locations therein.

Availability and reliability of information. Both sources of information, ecological and individual, can improve predictions. Assuming reliable assessment, relying on both predictors should normally inform better predictions than relying on only one predictor. That is, knowledge of $p(y \mid E_k, i)$ should be superior to knowledge of only $p(y \mid E_k)$. Thus, a specific location i nested in territory E_k should allow more accurate predictions than the overall ecology. However, to the extent that specific information about i is hard to assess, or low in reliability, due to its paucity or memory demands, relying on $p(y \mid E_k, i)$ may not be useful. If no specific information is available at all, or lost from memory, using $p(y \mid E_k, i)$ is impossible. Therefore, in a learning environment, in which individuating information is unreliable or difficult or even impossible to assess, ecological information $p(y \mid E_k)$ may lead to more accurate predictions than individuating information $p(y \mid E_k, i)$.

From baserate tendencies to PCs. Granting this rationale for baserate-driven response tendencies in one dimension, *y*, the algorithm can be easily extended to cover PCs between two dimensions, *y* and *x*. Thus, given skewed baserates in both *y* and *x* (e.g., a high predator rate and a high rate of dark places), the animal should acquire two simultaneous expectancies or response sets, for predicting many predators and for predicting many dark places in territory E_k. Especially when other ecologies predict fewer predators and fewer dark places, Y_{frequent} should be aligned with X_{frequent} and $Y_{\text{infrequent}}$ with $X_{\text{infrequent}}$ in any memory representation that is organized by categories or ecologies.

Even when the criterion rate Y_{frequent} does not vary from X_{frequent} to $X_{\text{infrequent}}$ (i.e., when the contingency is zero, or even reversed), the alignment of separately learned response tendencies for *Y* and *X* can be expected to produce a PC effect, for any of the following reasons: (a) The aligned association of frequent versus infrequent levels on both variables may induce a propositional inference that the two variables are correlated, which can then be used for guessing under uncertainty. This variant could be termed the syllogistic-reasoning account. (b) The more prevalent level of one variable, X_{frequent}, may activate or prime ecological information about E_k more effectively, prompting more reliable information on $p(y \mid E_k)$, than the less prevalent value, $X_{\text{infrequent}}$, which is less typical for E_k. This variant suggests why connectionist learning models (Dougherty, Gettys & Ogden, 1999; Fiedler, 1996) can easily account for PC inferences. (c) Given two uncorrelated but skewed distributions, the number of matching cases can still be larger than the number of non-matching cases.[7] If the cognitive concept of a contingency is represented in terms of an average match value (cf. White, 2001), this would imply that psychologically – though not statistically – PCs do imply contingencies.

Conclusion. To be sure, there is no guarantee that this idealized inference scheme will always lead to valid and rational decisions. However, it is exclusively based on normal rules of associative learning and adaptive response strategies, and the depicted processes through which baserate tendencies can be turned into PCs are not essentially flawed. If the assumption is met that ecological information is indeed more accessible and more reliable than individuating information, then there is hardly any viable alternative to PC inferences in an uncertain world in which predictions are indispensable.

The Informational Value of Baserate Information

According to this analysis, the adaptive value of PCs can be understood as a special case of a broader class of adaptive baserate strategies. Further evidence shows, indeed, that relying on criterion baserates alone, while ignoring contingencies with other variables, may inform correct predictions and decisions in many situations. Specifically, it has been shown that when the marginal baserates of a 2 x 2 table are sufficiently skewed, it may not be useful to utilize the contingency inherent in the diagonals; the marginals may be more informative than the diagonal (Goodie & Fantino, 1996, 1999; Kareev, 2000; Kareev et al., in press). For example, when a significant outcome occurs in 9 out of 10 cases under one condition and in 8 out of 10 cases under an alternative condition, then one should *not* conditionalize one's prediction on the condition, but rather predict the more frequent criterion value (with a baserate of 17 out of 20). Maximizing (i.e., predicting the most likely outcome) should be preferred, and actually is preferred to contingency-based strategies when the criterion distribution is sufficiently skewed (Kareev et al., 2008). Thus, the crucial condition that gives rise to PC effects, namely, skewed baserates, actually renders baserates informative and useful for predictions, regardless of whether an underlying contingency is understood or not

(Gaissmaier, Schooler & Rieskamp, 2006). Recent Monte-Carlo simulations by Freytag, Kutzner, Vogel, and Fiedler (2008a) indeed revealed a high accuracy of PC inferences under high skew. Viewed from this angle, PC inferences may be reflective of a general preparedness to exploit informative baserates.

Categorical coding. Another prominent topic of cognitive psychology that highlights the functionality of PCs and baserates lies in the formation of higher-order categorical memory codes. Encoding incoming information in terms of meaningful super-ordinate categories affords a highly economical and effective strategy, as it links new information to abstract knowledge structures in long-term memory (Huttenlocher, Hedges, & Vevea, 2000). In fact, PC-type representations can be considered a natural by-product of the formation of such economic memory codes, the advantages of which have been demonstrated in numerous experiments using categorized stimulus lists (Bousfield & Puff, 1965; Cohen, 1966; Fiedler, 1986; Tulving & Donaldson, 1972).

When a categorized list task calls for the processing of two attributes associated with multiple items, the formation and utilization of a categorical code is tantamount to forming PCs. A study by Fiedler and Graf (1990) illustrates this point, using countries as stimulus items and two attributes, the occurrence of a virus and the occurrence of a disease. In a first stimulus sequence, participants learned whether a virus was present or absent in 24 different countries, of which four belonged to each of six geographical categories. They formed a categorical code reflecting that the virus was present, say, in all Scandinavian, in part of the Mediterranean, and in no South-American countries. Then they learned, in a second run, whether a disease had occurred in the same countries. Again, a categorical code was formed of the differential occurrence rates of the disease in geographic categories. Not surprisingly, the actually existing correlation between virus and disease that was manipulated to be +.50 versus –.25 could not be assessed correctly when the matching and non-matching cases were flatly distributed across all six categories. When, however, high and low baserates of matching and non-matching cases were concentrated in different categories, offering a sound ecological explanation for the covariation of virus and disease, judgments were clearly sensitive to the sign of the correlation. Apparently, when there is a chance to form a categorical memory code, it is readily taken and utilized effectively.

In a study designed to manipulate category baserates independently of the actually existing correlation, Freytag, Vogel, Kutzner, and Fiedler (2008b, Exp. 2) presented their participants with a partner questionnaire filled in by both the male and female part in a couple seeking advice from a therapist. The questionnaire consisted of four subsets of dichotomous items in each of four domains (i.e., tenderness, debates, household, and sociability), representing self-descriptions of typical behaviors in a close relationship. Participants learned that in order to plan a therapy program, it was essential to assess the correlation between both partners' responses to the same items. In fact, *within* all domains, the (partial) correlation was negative (–0.33); that is, the proportion of Yes responses by the male partner was slightly higher for items the female partner had denied, and vice versa (see Table 4). However, the *ecological* correlation between the baserates of male and female Yes responses across the four domains was perfectly positive (+1.00), because both partners rejected most items in two domains and endorsed most items in the remaining two domains. The resulting total correlation across all 48 individual items was zero. Substantial PC effects in the range of +.25 to +.40 (pooling data across all four categories) were obtained in two dependent measures, memory reproductions of the male and female partners' responses to the original questionnaire items and predictions of their responses to new items of another, parallel test. Not surprisingly, high recall performance resulted from probability matching strategies (Erev & Barron, 2005; Shanks et al., 2002), that is, from simply reproducing the baserates of Yes responses per category. Subsequent explicit contingency judgments mirrored these guessing-based PC effects. These studies provide strong evidence for the functionality of prediction and recognition decisions that use the baserates of category-level memory representations. PC illusions can arise as a natural by-product of higher-order adaptive encoding strategies.

The adaptive advantage of baserates in general and PC strategies in particular was recently demonstrated in speeded-classification tasks, such as evaluative priming (Hermans, DeHouwer & Eelen, 1994) or the Implicit Association Test (IAT; Greenwald, Nosek & Banaji, 2003). In an IAT, implicit attitudes are inferred from the speed with which respondents sort two types of alternating stimuli, say, West or East German concepts, and positive or negative words, onto the same two response keys. West Germans' attitudes toward East Germans are supposed to be negative, when the latencies required to sort West / positive and East / negative onto the same keys are faster than the latencies needed to sort West / negative and East / positive together. In a recent study by Bluemke and Fiedler (2008), the baserates of West and East labels and of positive and negative stimuli were manipulated. The IAT results were affected by the stimulus baserates, independently of the attitudes to be measured. Both baserate-driven response biases (i.e., faster responding to more frequent stimulus classes) and PC effects were found. IAT effects were enhanced when the series included mostly West and mostly positive stimuli (or mostly East and mostly negative stimuli). Under this condition, the twofold response bias toward the more common categories facilitated sorting West / positive (and East / negative) together, because both response biases support the same response key. Opposing response biases reduced the speed of sorting West / negative and East / positive together. These findings not only suggest that strategies can obscure IAT measures. They also demonstrate neatly how motor response tendencies can contribute to adaptive PC strategies.

Interventions at Aggregate Level

So far, we have routed the adaptive value of PCs in the information value of category baserates and in the effectiveness of aggregate-level categorical memory codes. Contingency assessment, however, is not only for accurate impression formation and judgment, it is also essential for goal-directed behavior in the service of an organism's needs and preferences.

The pragmatic value of the PC *illusion* should be evident whenever ecological correlations are more useful and more functional than individuating correlations. Some reflection reveals that this may be often the case when it comes to assessing the impact of interventions and "natural experiments." Just as experimental manipulations apply to aggregates or groups of participants, many real-life interventions are only feasible at the aggregate level of ecologies rather than individual targets. A teacher can change his or her teaching strategy, or tone, or the work load, but the effects of such interventions is only visible at the level of the class, or ecology; it is rarely possible to achieve differential impact in different individual students. Likewise, an animal striving to survive in an uncertain world can avoid certain ecologies, or approach other areas, or make behavior contingent on certain time periods or settings, but it can hardly learn to avoid individual predators or to seek a specific food item.

If our goal is to understand the contingencies that hold between causes and effects, or signals and significant stimuli, then it may be more prudent to apply a nomothetic strategy, trying to control ecologies rather than individual cases, just as it is often easier to obtain scientific evidence at group level than at individual level. To the extent that interventions and attempts to control an uncertain environment do not come with the guarantee of a perfectly predictable outcome, but instead can only be assumed to affect the baserate of the significant events to be controlled, a functional assessment of causality must be sensitive to ecological baserates rather than individual events. Assessing and controlling ecologies may thus be a more feasible strategy than trying to gain control over individual events.

PCs in the Researcher's Mind

What holds for active interventions in real life, should apply to experimental manipulations and research designs as well. It is therefore not too surprising that PCs have also intruded into the one domain in which rationality and logic are weighted most strongly, namely, in science. Some of the methods that have become common practice and some of the theories that seem firmly entrenched in our discipline's ground knowledge are based on PC inferences. A glance at the role PCs play in science should help substantiate the conclusion that the apparent illusion cannot be discarded as simply irrational.

PC-type interpretations of empirical results. Consider the fact that most behavioral research uses group data to investigate individual cognitive or affective processes. Closer inspection reveals that relying on group data produces PCs in both quasi-experimental and experimental studies. The most obvious instantiation of PCs can be found in studies pitting extreme groups against each other. When juxtaposing, say, the 10% most successful students and the 10% least successful students, a common mistake is to believe that significant differences of these extreme groups on other variables (e.g., socio-economic status, SES) account for the contrast between these extremely different groups. Even when the top group ranges much higher in SES than the bottom group, this need not imply that SES correlates with performance of individual students.

What is less obvious, though, is that even experiments using randomized groups may not allow for inferences beyond the PC level. Thus, if TV consumption is manipulated between two groups and subsequent aggression is higher in the high-TV group than in the low-TV group, all we know is that the average TV exposure correlates with the average aggression per group. Typical manipulation checks refer to group means; so the TV treatment could have had quite different influences on individual school children. Any causal interpretation of group data in terms of individual-level theoretical constructs – such as the child's identification, excitation transfer, or catharsis – is therefore based on a PC inference from group baserates to individual processes.

Pseudocontingent explanations. The least obvious form of PC reasoning in science can be found in highly familiar theoretical arguments, the constituent elements of which are so firmly entrenched in our background knowledge that their problematic character easily evades the attention of researchers and theorists. Evolutionary accounts of behavioral phenomena, for instance, rely on the assumption that long-term correlations supposed to hold at the aggregate level of thousands of generations can explain short-term correlations within a single generation, or even a single experiment.

One prominent example of PC reasoning was uncovered in Nickerson's (1995) critique of Kahneman and Knetsch's (1992) finding that the median of all people's willingness to pay for different public goods was highly correlated with the average person's moral satisfaction associated with these goods. The theoretical interpretation that this correlation of group averages supports the hypothesis that moral satisfaction drives individual persons' willingness to pay constitutes a PC inference, because within-person correlations need not resemble the correlation of averages.

Another notable example reveals a PC-argument in the explanation of confirmation biases in social hypothesis testing. Zuckerman, Knee, Hodgins, and Miyake (1995) noted that the tendency of interviewers to ask one-sided questions is not sufficient to explain the confirmation bias. Asking mostly interview questions about extraverted behaviors is not sufficient to explain that interviewers believe to have confirmed that the interviewee is extraverted. To make the confirmation complete, it was argued that interviewers' questioning bias has to be complemented by an acquiescence bias – the tendency of interviewees to provide predominantly confirming responses. Thus, a high baserate of extraversion questions together with a high baserate of confirming responses was considered sufficient to prove that the interviewee is extraverted. Logically, such an inference is unwarranted, because the interviewee's confirmation rate for the minority of introversion questions may be even higher than the confirmation rate for the majority of extraversion questions. Two skewed baserate distributions *do not* make for a contingency. Researchers, reviewers, and even opponents hardly notice the PC inference that underlies many theoretical arguments. Theoretical derivations based on two or more well-established tendencies may render theoretical arguments pseudocontingent – a typical by-product of nomothetic research programs.

Conclusions

Summary: Generality of the Phenomenon

To reiterate the major theoretical and empirical message to be conveyed in this article, we have reviewed a growing body of evidence showing that when organisms seem to assess and utilize contingencies, they may in fact process pseudocontingencies (PCs). Rather than being sensitive to how the values of one variable change with the values of another, contingency judgments, predictions, and evaluations often reflect the alignment of the two variables' separate baserate distributions. Speaking in technical but familiar terms, subjectively inferred contingencies may follow the marginals rather than the entries of contingency tables. If the baserates of two variables tend to correlate positively (negative) across ecologies, or if the baserates observed in a single ecology are skewed in the same (opposite) direction, the correlation between individual cases within ecologies will also be inferred to be positive (negative). If there are no correlated trends in ecological baserates, no PC inference will be made.

A review of pertinent studies revealed systematic and often quite strong PC illusions obtained across various stimulus materials and task conditions. PCs were demonstrated in such diverse content domains as social stereotypes, group discrimination, student achievement judgments, assessment of personality traits, diagnostic and economic judgments, and memory for questionnaire responses. With respect to task conditions, PCs were shown for dichotomous attributes, for three variable levels as well as for continuous variables. The illusion appeared to be strongest when two contrasting ecologies or groups have divergent baserates for two variables, yielding a perfect ecological correlation that highlights the skewness of the variables within each of the ecologies.

However, PCs also generalized to the case of a single ecology, in which an ecological correlation must be inferred implicitly, or to multiple ecologies with less clear-cut ecological correlations. A single, unitary principle seems to account for all the variants of the illusion: PCs originate in a learning environment that makes ecological correlations (at a higher level of aggregation) more accessible, more manipulable, and more useful than individuating correlations (at a lower level of aggregation). No extra account is required for single-ecology PCs. Although the ecological correlation cannot be observed directly in this case, it can be inferred indirectly from a comparison of observed baserate trends in the given ecology with default expectancies.

The distinct nature of PCs as opposed to genuine contingencies is most clearly demonstrated when information about two variables is presented successively, in separate runs that prevent the encoding of joint observations of paired variable values. However, PCs also intrude into the domain of genuine contingencies, when pair-wise values on both variables are presented simultaneously. When the contingency implied by the joint observations diverges from the PC implied by the baserates, the latter will often override the former, suggesting a process that is more sensitive to PCs than to contingencies proper.

The persistent finding of PC illusions was obtained across different causal models induced to make sense of ecological correlations and regardless of the explicit instruction to attend to higher or lower aggregation levels, although the size of the illusion seems to be moderated by such boundary conditions. PCs arose when baserates were presented propositionally, using verbal quantifiers or numerical statistics to describe skewed attribute baserates. We have pointed out the structural similarity of PC-type inferences to atmospheric biases in syllogistic reasoning: Mostly X in Z and mostly Y in Z seems to imply a positive statistical relationship between X and Y. Given this analogy to deductive reasoning, it is no surprise that PC arguments can also be found in scientists' theoretical reasoning.

PCs and Different Causal Models

Thus, while we have repeatedly emphasized the distinct nature of the cognitive process underlying PC inferences, it is equally important to point out that PCs are not bound to one specific causal model for the correlated variables *X*, *Y* and the ecological variable E. The crucial feature of PC inferences consists in the carry-over of statistical trends that hold at a high level of ecological baserates to predictions of specific events at a lower level of aggregation. Just as a correlation between X and Y is compatible with many different causal models, PC effects are not confined to one specific causal process.

To illustrate this important point with a final example, a teacher's firm belief that sanctions (*X*) increase motivation (*Y*) may reflect an apparent correlation between the means or baserates of sanctions and high motivation across ecologies. Even though the motivation for specific behaviors within ecologies may not correlate or even correlate negatively with the use of sanctions, a PC effect lets the teacher believe that sanctions increase motivation on specific tasks. Ecologies could be different types of students, subject matters, or time periods.

The crucial condition here is the existence and effective encoding of an ecological correlation. However, the PC inference is independent of the causal origin of the ecological correlation. One possibility is that ecologies (E) mediate a causal influence of sanctions (*X*) on motivation (*Y*). This causal model, $X \rightarrow$ E $\rightarrow$ Y, would be appropriate when an increased rate of sanctions creates an ecological condition (e.g., an affective atmosphere) which in turn enhances students' motivation. Conversely, E could mediate the impact of the general motivation level on the rate of sanctions, $Y \rightarrow$ E $\rightarrow$ X, for instance, when the students' enhanced motivation creates "loaded situations" that motivate the teacher to increase his or her rate of sanctions. A third possibility is that ecologies are a common cause of trends in sanctioning and in motivation, E $\rightarrow$ X, Y. Thus, highly significant or existential situations may increase the rate of both behaviors, which may otherwise be uncorrelated. Still another possibility is that an existing correlation between *X* and *Y* may be moderated by E.

There is no logical or psychological reason to restrict PCs to specific causal origins of ecological trends. The empirical findings we have reviewed testify to PCs of different causal types, as directly shown by Fiedler and Freytag (2004). Theoretically, then, PCs must not be conceived as re-interpretations of specific causal relations but as a general inference scheme that apply to ecological trends of any origin. This is not to preclude, though, that empirical research will show PC illusions to be more likely in some causal constellations than in others (consistent with the tentative finding of somewhat stronger PCs for common cause than common effect models; Fiedler & Freytag, 2004).

Limitations and Boundary Conditions

Despite the generality of PC inferences, some findings and theoretical arguments suggest that there must also be limits to the phenomenon. For several reasons, indeed, PC effects may be undone or even reversed under specific conditions.

A-priori constraints. One major limitation stems from the inherent meaning of the attributes whose contingency is to be assessed. The mere co-occurrence of baserate trends may not be a sufficient condition for a PC effect. Not each and every pair of variables with jointly skewed baserates must produce a PC. For example, everyday traffic experience shows that most vehicles have four wheels and most streets have a grey surface, but we would hardly form an illusory association between grey and four wheels. Just as a-priori constraints restrict causal inference (called "causal power", cf. White, 1995) and even conditioning (called "preparedness"; Öhman, 2001), there should be top-down constraints on PC inferences too. Some attribute pairs are more likely to be related, when their baserates are jointly skewed, than others. We would suggest the term integral to denote attributes that are more likely to be interrelated in a PC than other, less integral attributes (for a discussion of integrality, see Garner, 1970, or Maddox, 1992). An open empirical and theoretical question is: What determines the propensity to form a mentally represented PC between two variables given appropriate baserate data?

A-posteriori constraints. The second limiting condition is a-posteriori in nature. The PC-type dominance of high-aggregate information over low-aggregate individuating data is contingent on the pragmatic task context. Decision makers or organisms more generally, do have the capacity to learn and utilize low-level, individuating information when supported by the reinforcement structure of the task. A prominent PC counter-example, in which individuating observations dominate aggregate baserates, can be found in multi-trial dilemma games (Dawes, 1980). The reinforcement structure of such experimental games typically favors the encoding of single-trial outcomes, rather than aggregate outcomes across the entire series (cf. Fiedler, 2008). As the payoff is presented immediately after every trial, it is no wonder that players are most sensitive to information at the level of singular trials. That is, players are likely to defect rather than to cooperate because payoffs are higher for defection than for cooperation on every single trial. In the long run, though, cooperation would yield more success than defection, but the immediate payoff prevents players from forming aggregate records across longer sequences of trials. Would the payoff structure support such aggregate coding – for instance, by providing outcome feedback after blocks of 20 trials – the cooperation advantage might be discovered more readily.

Dilemma games highlight that aggregate-level information need not always be encoded more efficiently than low-level information (Fiedler, 2008; Vlaev & Chater, 2006). The same conclusion applies more generally to delay of gratification problems (Metcalfe & Mischel, 1999), in which the short term advantages of one behavior loom larger than the long-term advantages of alternative behaviors. Another prominent example can be found in the shared information effect in group-decision making (Mojzisch & Schulz-Hardt, 2006; Stasser & Titus, 1985), showing that individual group members' preferences override the aggregate, group-level preferences. A common denominator of these counter-examples to the PC rule is that both payoff schemes and encoding conditions favor information at low levels of aggregation.

PCs and exemplar-based processes and representations. With regard to cognitive representations modes, PC inferences are clearly at variance with the assumptions of exemplar-memory models (Juslin & Persson, 2002; Nosofsky & Johansen, 2000). Whereas these models assume that individuating stimulus traces are stored in memory to determine subsequent judgments and decisions, PCs rely on the opposite assumption that individual exemplars are often hard to encode, not available, or easily forgotten. Both types of information representation have their place in adaptive cognition, and we indeed believe that the domain boundary of exemplar models describes one of the major limitations of the PC domain. An intriguing suggestion for future research, then, is that successfully established exemplar memories should offer a remedy against PC illusions.

However, even when future research reveals many ways in which a-priori knowledge, stimulus representation formats, and pay-off structures moderate PC effects, this would be in the spirit of the basic idea underlying the PC approach that there is more than a single cognitive algorithm for assessing contingencies. When organisms make inferences between two variables, they are not bound to one procedure preferred in statistics textbooks. They do not always look at joint frequencies. Instead, they may utilize baserate trends, especially when baserates are skewed and informative. Outside statistical information, there are still other cues that organisms can utilize, such as spatial-temporal contiguity, or the similarity or intensional relatedness of observations (Fiedler, 2000; Fiedler & Ströhm, 1986).

Setting PCs Apart from Related Paradigms and Conceptions

Quite distinct from those factors that limit or override the occurrence of PC inferences when the theoretical conditions for PCs are met, other concepts and paradigms are fundamentally different and should thus be clearly distinguished from PCs.

Contingencies and causal inference. The PC paradigm does not conflict with the huge research program on causal learning, conditioning, and probabilistic inference, which testifies to organism's sensitivity to genuine contingencies (Allan, 1993; Fiedler, 2000). Both inference schemes complement rather than contradict each other. Certain conditions (i.e., the availability and encodability of joint observations and non-skewed distribution of a criterion variable) facilitate the learning and utilization of genuine contingencies, whereas other conditions (isolated assessment of correlated variables, skewed distributions) only allow for PC-type inferences. Several recent approaches converge in the contention that it can be more useful to base inferences and predictions on baserates rather than contingencies proper (Allan et al., 2008; Erev & Barron, 2005; Gaissmaier et al., 2006). The findings reviewed in the present article suggest that the domain of PC inferences may be larger than expected and sometimes even intrude into the domain of contingency studies. Future research must clarify the division of labor between both camps.

Simpson's paradox. The trivariate context of PC inferences, with ecologies affording a third variable in addition to the two focal variables to be correlated, is reminiscent of other trivariate paradigms that resemble PCs in some, but not in all aspects. Most prominently, Simpson's (1951) paradox constitutes a psychological analog of a spurious correlation, as described in the empirical review above. For instance, judges are puzzled when across all universities female applicants to graduate programs are more often rejected than males, even though females are more successful than males within individual universities. Judges fail to understand that the partial correlation between gender and success (within universities) can be opposite in sign to the zero-order total correlation. This is possible when universities with high rejection rates have high baserates of female applicants, whereas universities with low rejection rates have low female baserates. However, in Simpson's paradox, partialling out the third variable always accounts for (at least part of) the zero-order correlation so that the role of the ecological variable is confined to demonstrating that the correlation is spurious. This restriction does not hold for PCs, which allow for all kinds of combination of correlations within and between categories or ecologies. PCs even arise in singular ecologies, when the third variable to be partialled out is undefined.

Suppressor effect. Suppressor effects (Conger & Jackson, 1972) can be considered the opposite of Simpson's paradox, or spurious correlations, in that the third variable does not account for, or explain away, the zero-order correlation between two focal variables. Rather, partialling out the suppressor renders the correlation stronger, as the suppressor accounts for error variance associated with the two focal variables. Again, this is but a special case of the conditions leading to PC inferences. Like spurious correlations, suppressor relations are undefined when there is but one ecology, when one of the three pair-wise correlations is unknown, or when no multivariate observations exist.

Discounting effects. An intriguing theoretical implication is that ecological correlations resist discounting effects in attribution (Kelley, 1973; McClure, 1998; Morris & Larrick, 1995). Thus, when both partners' questionnaire responses in the Freytag et al. (2008b) study tend to be positive in two domains but negative in two other domains, one could attribute this category-level correlation to differences between domains and discount the agreement between partners as an underlying cause. Maybe, the tenderness domain elicits mostly positive responses, and the housework domain mostly negative responses, from everybody, not just this couple. Such a discounting was not observed, though. When two further respondents' consensual endorsement and denial of the same questionnaire domains highlighted an external, domain-specific attribution, judges did not discount a partner-specific cause. Rather, they readily inferred high agreement between partners on specific items. The co-presence of an ecological cause reinforced rather than discounting the inference of an individual-level cause (see also Fiedler, Walther, & Nickel, 1999; Trope, 1986).

Adaptive Cognition in Probabilistic Environments

Throughout this article, we have treated PCs not merely as irrational illusions, but also as cognitive devices that can serve adaptive functions in a complex, probabilistic world. As the environment rarely provides us with multivariate observations but most often only with partial information about singular attributes at a time, organisms have to resort to inference schemes different from genuine contingencies. Forming PCs from systematic baserate differences observed at the aggregate level of ecologies, categories, or groups may be the best that organisms can do in this situation. Most of the time, when ecological correlations between baserates point in the same direction as individuating correlations at lower levels, PC inferences will lead to useful, satisficing predictions and decisions (Simon, 1982).[8] Like all illusions, they only lead to errors, faulty decisions, and costs, when generalized to new contexts within which the illusion is dysfunctional.

Nevertheless, because such problem contexts do exist, it has to be acknowledged that PC illusions can be blatantly wrong, informing seriously flawed judgments and decisions. As typical of all illusions, when there is no fit between the logic underlying PC inferences and the structure of the problem environment, the illusion can produce seriously wrong decisions, with unjust and expensive consequences. When PCs distort the evaluation of individual students' performance toward the class baserates (Fiedler et al., 2007), the outcome is unfair, however adaptive and rational PC inferences may be in other situations. Similarly, when group data are used to validate intrinsically individual processes (Nickerson, 1995), the PC outcome is irrational. Or, when group-level trends are applied to individual group members (Hoffman & Hurst, 1990), the PC inference is at the heart of social stereotypes.

The conditions under which PCs produce erroneous and potentially harmful consequences have been stated repeatedly. If the *ecological* correlation $r({}_{\text{mean}}X_k, {}_{\text{mean}}Y_k)$ diverges from the *partial* correlation $r(x_i,y_i|E)$ and hence from the *total* individual correlation $r(x_i,y_i)$, PC inferences will likely lead to errors when two additional consequences are met: (a) the problem structure calls for judgments of the relation between individual attributes x_i and y_i; but (b) the task context favors the assessment of ecological baserates ${}_{\text{mean}}X_k$ and ${}_{\text{mean}}Y_k$.

In these situations, it makes little sense to justify PC-driven category mistakes with the argument that PCs are consistent with Bayesian logic. Although this argument is potentially correct, it needs to be specified. Given elevated baserates of high X and Y values, the posterior distribution of possible $r(x_i,y_i)$ correlations shows that there are more possibilities for x and y to correlate positively and the range of negative correlations is restricted ($r > -1$). This Bayesian argument holds when all possibilities of combining data from the two distributions are equally likely. However, being in a world in which the ecological correlation diverges strongly from the pooled partial correlation implies that, within specific ecologies, some possibilities are much more or much less likely than in other ecologies.

We therefore return to the question of what "priors" should be used for a Bayesian solution, the priors that hold across all ecologies or the specific priors within specific ecologies. Given the divergence between partial correlation and ecological correlation, the two sets of priors should differ markedly. Bayesian calculus itself does not solve this problem, and it cannot provide a general absolution for all mistakes made in PC's name. Whether a problem has to be tackled at a high or low level of aggregation cannot be answered on Bayesian ground. The answer can only be found in a functional analysis of the causal and hedonic structure of the problem environment.

References

Allan, L.G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin, 114*, 435-448.

Allan, L.G., Hannah, S.D., Crump, M.J.C., & Siegel, S. (2008). The psychophysics of contingency assessment. *Journal of Experimental Psychology: General, 137*, 226-243.

Alloy, L.B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: the joint influence of prior expectations and current situational information. *Psychological Review, 91*, 112-149.

Arieh, Y., & Algom, D. (2002). Processing picture – word stimuli: The contingent nature of picture and of word superiority. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 221-232.

Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review, 97*, 548-564.

Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review, 6*, 57-86.

Bluemke, M., & Fiedler, K. (2007). *Base rate effects on the IAT.* Manuscript submitted for publication.

Bousfield, W.A., & Puff, C.R. (1965). Determinants of the clustering of taxonomically and associatively related word pairs. *Journal of General Psychology, 73*, 1965, 211-221.

Chater, N., & Oaksford, M. (Eds.), *The probabilistic mind:Prospects for rational models of cognition.* London: Oxford University Press.

Chater, N., Vlaev, I., & Grinberg., M. (2007). *A new consequence of Simpson's paradox: Stable cooperation in single-shot Prisoner's Dilemma from populations of individual learners.* Manuscript submitted for publication.

Cohen, B.H. (1966). Some-or-none characteristics of coding behavior. *Journal of Verbal Learning and Verbal Behavior, 5*, 182-187.

Conger, A. J., & Jackson. D. N. (1972). Suppressor variables, prediction, and the interpretation of psychological relationships. *Educational and Psychological Measurement, 32.* 579-599.

Crocker, J. (1981). Judgment of covariation by social perceivers. *Psychological Bulletin, 90*, 272-292.

Dawes, R. M. (1980). Social dilemmas. *Annual Review of Psychology, 31*, 169–193.

DeHouwer, J. & Beckers, T. (2002). A review of recent developments in research and theories on human contingency learning. *Quarterly Journal of Experimental Psychology, 55*, 289-310.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review, 106*, 180-209.

Eder, A.B., Fiedler, K., & Hamm-Eder, S. (2008). *Illusory correlations revisited: The role of pseudocontingencies and working-memory capacity*. Manuscript submitted for publication.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review, 112*, 912-932.

Fairchild, H. H. (1991). Scientific racism: The cloak of objectivity. *Journal of Social Issues, 47*, 101-115.

Fiedler, K. (1986). Person memory and person judgments based on categorically organized information. *Acta Psychologica, 1986*, 61, 117-135.

Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review, 103*, 193–214.

Fiedler, K. (2000). Illusory correlations: A simple associative algorithm provides a convergent account of seemingly divergent paradigms. *Review of General Psychology, 4*, 25-58.

Fiedler, K. (2008). Pseudocontingencies – A key paradigm for understanding adaptive cognition. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind:Prospects for rational models of cognition.* London: Oxford University Press.

Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology, 87*, 453-467.

Fiedler, K., Freytag, P., & Unkelbach, C. (2007). Pseudocontingencies in a simulated classroom. *Journal of Personality and Social Psychology, 92*, 665-677.

Fiedler, K., & Graf, R. Grouping and categorization in judgment of contingency. In J.P. Caverni, J.M. Fabre, & M. Gonzalez (Eds.), *Cognitive biases*. North-Holland, 1990.

Fiedler, K., Russer, S., & Gramm, K. (1993). Illusory correlations and memory performance. *Journal of Experimental Social Psychology, 29*, 111-136.

Fiedler, K., & Stroehm, W. (1986). The use of statistical, spatial-temporal, and intensional information in judgments of contingency. *European Journal of Social Psychology, 16*, 385-398.

Fiedler, K., Walther, E., Freytag, P., & Plessner, H. (2002). Judgment biases in a simulated classroom--A cognitive-environmental approach. *Organizational Behavior and Human Decision Processes, 88*, 527-561.

Fiedler, K., Walther, E., & Nickel, S. (1999). Covariation-based attribution: On the ability to assess multiple covariates of an effect. *Personality and Social Psychology Bulletin, 25*, 607-622.

Freytag, P. (2003). *Contextually determined typicality*. Berlin: Logos.

Freytag, P., Kutzner, F., Vogel, T., & Fiedler, K. (2008a). On the adaptive value of pseudocontingencies. Unpublished research, University of Heidelberg.

Freytag, P., Vogel, T., Kutzner, F., & Fiedler, K. (2008b). Pseudocontingencies due to base-rate reproduction. Unpublished research, University of Heidelberg.

Gaissmaier, W., Schooler, L.J., & Rieskamp, J. (2006). Simple predictions fueled by capacity limitations: When are they successful? *Journal of Experimental Psychology: Learning Memory and Cognition, 32*, 966-982.

Garner, W. R. (1970). The stimulus in information processing. *American Psychologist, 25*, 350-358.

Goldstein, H. (1999). *Multilevel statistical models*. London: Arnold. Book Citation.

Goodie, A. S., & Fantino E. (1996). Learning to commit or avoid the base-rate error. *Nature*, 380, 247-249.

Goodie, A. S., & Fantino E. (1999). What does and does not alleviate base-rate neglect under direct experience. *Journal of Behavioral Decision Making*, 12, 307-335.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology, 85*, 197-216.

Hamilton, D.L. (1981). Illusory correlations as a basis for stereotyping. In D.L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior*. Hillsdale, N.J.: Erlbaum.

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12*, 392-407.

Hammond, J.L. (1973). Two sources of error in ecological correlations. *American Sociological Review, 38*, 764-777.

Hermans, D., DeHouwer, J., & Eelen, P. (1994). The affective priming effect: Automatic activation of evaluative information in memory. *Cognition and Emotion, 8*, 515-533.

Hoffman, C. & Hurst, N. (1990). Gender stereotypes: Perception or rationalization? *Journal of Personality and Social Pyschology, 58*, 197-208.

Huttenlocher, J. & Hedges, L., & Vevea, J.L. (2000). Why Do Categories Affect Stimulus Judgement? *Journal of Experimental Psychology: General, 129*, 220-241.

Juslin, P., & Persson, M. (2002). PROBabilities from Exemplars (PROBEX): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science, 26*, 563–607.

Kahneman, D., & Knetsch, J.L. (1992). Valuing public goods: The purchase of moral satisfaction. *Journal of Environmental Economics and Management, 22*, 57-70.

Kareev, Y. (2000). Seven (indeed, plus minus two) and the detection of correlations. *Psychological Review, 107,* 397–402.

Kareev, Y., Fiedler, K. & Avrahami, J. (in press). Base rates, contingencies, and prediction behavior. *Journal of Experimental Psychology: Learning, Memory & Cognition.*

Kelley, H.H. (1973). The process of causal attribution. *American Psychologist, 28*, 107-128.

Klauer, K. C., & Meiser, T. (2000). A source-monitoring analysis of illusory correlations. *Personality and Social Psychology Bulletin, 26*, 1074-1093.

Klauer, K.C., Musch, J. & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review, 107*, 852-884.

Kutzner, F., Freytag, P., Vogel, T., & Fiedler, K. (in press). Base-rate neglect as a function of base rates in probabilistic contingency learning. *Journal for the Experimental Analysis of Behavior.*

Lambert, A.J., & Wyer, R.S. (1990). Stereotypes and social judgment: The effects of typicality and group heterogeneity. *Journal of Personality and Social Psychology, 59*, 676-691.

Maddox, W. T. (1992). Perceptual and decisional separability. In F. G. Ashby (Ed.), *Probabilistic multidimensional models of perception and cognition* (pp. 147-180). Hillsdale, NJ: Erlbaum.

Marsh, H. W. & Hau, K. T. (2003). Big fish little pond effect on academic self-concept: A cross-cultural (26 country) test of the negative effects of academically selective schools. *American Psychologist, 58*, 364-376.

McClure, J. (1998). Discounting of causes of behavior: Are two reasons better than one? *Journal of Personality & Social Psychology, 74*, 7-20.

McGarty, C., Haslam, S. A., Turner, J. C., & Oakes, P. J. (1993). Illusory correlation as accentuation of actual intercategory difference: Evidence for the effect with minimal stimulus information. *European Journal of Social Psychology, 23*, 391-410.

Meiser, T. (2003). Effects of processing strategy on episodic memory and contingency learning in group stereotype formation. *Social Cognition, 21*, 121-156.

Meiser, T. (2006). Contingency learning and biased group impressions. In K. Fiedler & P. Juslin (eds.), *Information sampling and adaptive cognition* (pp. 183-209). Cambridge: Cambridge University Press.

Meiser, T., & Bröder, A. (2002). Memory for multidimensional source information. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 116-137.

Meiser, T., & Hewstone, M. (2001). Crossed categorization effects on the formation of illusory correlations. *European Journal of Social Psychology, 31*, 443-466.

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*, 599-614.

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*, 315-336.

Metcalfe, J., & Mischel, W. (1999). A hot-cool analysis of the delay of gratification: the dynamics of willpower. *Psychological Review, 106*, 3-19.

Mojzisch, A., & Schulz-Hardt, S. (2006). Information sampling in group decision making. In K. Fiedler & P. Juslin (Eds.), *Sampling and adaptive cognition* (pp. 299-326). New York: Cambridge University Press.

Morris, M. W., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of Discounting in causal attribution. *Psychological Review,* 102, 331-355.

Nickerson, C.E. (1995). Does willingness to pay reflect the purchase of moral satisfaction? A reconsideration of Kahneman and Knetsch. *Journal of Environmental Economics and Management, 28*, 126-133.

Nosofsky, R. M., & Johansen, M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review, 7*, 375–402.

Öhman, A. (2001). Fears, phobias, and preparedness: Toward an evolved model of fear and fear conditioning. *Psychological Review, 108*, 483-522.

Pedhazur, E.J. (1982). *Multiple regression in behavioral research* (2nd ed.). New York: Holt, Rinehart, and Winston.

Reips, U.-D., & Waldmann, M. R. (2008). When learning order affects sensitivity to base rates. Challenges for theories of causal learning. *Experimental Psychology, 55*, 9-22.

Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review, 15*, 351-357.

Schaller, M., Asp, C. H., Rosell, M. C., & Heim, S. J. (1996). Training in statistical reasoning inhibits the formation of erroneous group stereotypes. *Personality and Social Psychology Bulletin, 22*, 829-844.

Schaller, M., & O'Brien, M. (1992). "Intuitive analysis of covariance" and group stereotype formation. *Personality and Social Psychology Bulletin, 18*, 776-785.

Shanks, D. R., Tunney R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making, 15*, 233-250.

Simon, H.A. (1982). *Models of bounded similarity*. Cambridge, MA: MIT Press.

Simpson, E. H. (1951). The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society, Series B, 13*, 238-241.

Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling in group discussion. *Journal of Personality and Social Psychology, 48*, 1467-1478.

Stroessner, S. J., Hamilton, D. L., & Mackie, D. M. (1992). Affect and stereotyping: The effect of induced mood on distinctiveness-based illusory correlations. *Journal of Personality and Social Psychology, 62*, 564-576.

Trope, Y. (1986). Identification and inference processes in dispositional attribution. *Psychological Review,* 93, 239-257.

Tulving, E., & Donaldson, W. (1972). *Organization of memory.* Oxford, England: Academic Press.

Vlaev, I., & Chater, N. (2006). Game Relativity: How Context Influences Strategic Decision Making. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32*, 131-149.

Waldmann, M., & Hagmayer (1995). Causal paradox: When a cause simultaneously produces and prevent an effect. *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

White, P.A. (1995). Use of prior beliefs in the assignment of causal roles: Causal powers versus regularity-based accounts. *Memory & Cognition, 23*, 243-254.

White, P.A. (2001). Causal judgment from contingency information: Relation between subjective reports and individual tendencies in judgment. *Memory & Cognition, 28*, 415-426.

Woodworth, R.S., & Sells, S.B. (1935). An atmosphere effect in formal syllogistic reasoning. *Journal of Experimental Psychology, 18*, 451- 460.

Zuckerman, M., Knee, C. R., Hodgins, H. S. & Miyake, K. (1995). Hypothesis confirmation: The joint effect of positive test strategy and acquiescence response set. *Journal of Personality and Social Psychology, 68*, 52-60.

Table 1. Attitudes regarding Freedom of Science (F) and Humanism (M) of 16 students in the simulated classroom used to induce positive (PC+) and negative (PC–) Pseudocontingencies.

	Positive PC		Negative PC	
Subgroup	F Attitude	H Attitude	F Attitude	H Attitude
Blue				
Student #1	Pro	Contra	Pro	Pro
Student #2	Pro	Contra	Pro	Pro
Student #3	Pro	Pro	Pro	Contra
Student #4	Pro	Pro	Pro	Contra
Student #5	Neutral	Pro	Neutral	Contra
Student #6	Neutral	Pro	Neutral	Contra
Student #7	Neutral	Pro	Neutral	Contra
Student #8	Neutral	Pro	Neutral	Contra
Red				
Student #1	Neutral	Contra	Neutral	Pro
Student #2	Neutral	Contra	Neutral	Pro
Student #3	Neutral	Contra	Neutral	Pro
Student #4	Neutral	Contra	Neutral	Pro
Student #5	Contra	Contra	Contra	Pro
Student #6	Contra	Contra	Contra	Pro
Student #7	Contra	Pro	Contra	Contra
Student #8	Contra	Pro	Contra	Contra

Note. Only H attitudes varied as a function of the sign of the resulting pseudocontingency.

Table 2. Frequency distributions used in Fiedler and Freytag (2004).

a) Exp.1

PC condition	Ward A Symptoms Low	Ward A Symptoms High		Ward B Symptoms Low	Ward B Symptoms High		Total Symptoms Low	Total Symptoms High	
Positive									
Diet									
Vegetarian	(3)	(9)	12	(27)	(9)	36	(30)	(18)	48
Prebiotic	(9)	(27)	36	(9)	(3)	12	(18)	(30)	48
	12	36		36	12		48	48	
Zero									
Diet									
Vegetarian	(12)	(12)	24	(12)	(12)	24	(24)	(24)	48
Prebiotic	(12)	(12)	24	(12)	(12)	24	(24)	(24)	48
	24	24		24	24		48	48	
Negative									
Diet									
Vegetarian	(9)	(3)	36	(9)	(27)	12	(30)	(18)	48
Prebiotic	(27)	(9)	12	(3)	(9)	36	(18)	(30)	48
	12	36		36	12		48	48	

b) Exp.2 / Exp.3

Test X	Group A Test Y Low	Group A Test Y High	Group B Test Y Low	Group B Test Y High	Total Test Y Low	Total Test Y High
Low	4	8	16	8	20	16
High	8	16	8	4	16	20

c) Exp.3

Test X	Group A Test Y Low	Group A Test Y Inter	Group A Test Y High	Group B Test Y Low	Group B Test Y Inter	Group B Test Y High	Total Test Y Low	Total Test Y Inter	Total Test Y High
Low	0	0	6	12	6	6	12	6	12
Intermediate	0	0	6	6	0	0	6	0	6
High	6	6	12	6	0	0	12	6	12

Table 3. Mean predicted values obtained by Fiedler and Freytag (2004, Exp. 1) as a function of target group (Ward A vs. Ward B), feature (symptom checklist *vs.* diet), level of given value (low *vs.* high checklist score or vegetarian *vs.* prebiotic diet, respectively), and PC type (positive *vs.* zero *vs.* negative).

	Ward A				Ward B			
	Checklist		Diet		Checklist		Diet	
PC Type	Low	High	VEG	PRE	Low	High	VEG	PRE
Positive	0.25	0.86	0.34	0.66	0.19	0.75	0.38	0.70
Zero	0.42	0.42	0.52	0.50	0.44	0.44	0.46	0.52
Negative	0.78	0.19	0.64	0.36	0.72	0.38	0.66	0.25

Note. VEG = vegetarian diet, PRE = prebiotic diet. All dependent measures were linearly transformed to the range of 0 (minimum) to 1 (maximum).

Table 4. Frequency distributions used in Freytag, Vogel, Kutzner, and Fiedler (2008, Exp. 2).

	Domain 1 Male Partner		Domain 2 Male Partner		Domain 3 Male Partner		Domain 4 Male Partner		Total Male Partner	
Female partner	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
No	6	3	0	3	6	3	0	3	12	12
Yes	3	0	3	6	3	0	3	6	12	12

Footnotes

[1] There may be more than two aggregation levels in a generally formulated model.

[2] Note that the ecological correlation is always perfect if only two ecologies with differing baserates of X and Y are concerned, because two points in a plane can be perfectly fitted by one line.

[3] Although marginal distributions that are skewed in the same direction can result from tables with negative correlations, and although oppositely skewed marginal distributions can result from tables with positive correlations, skewed baserates restrict the range of the correlation. Baserates skewed in the same direction imply that the lower bound of the correlation is larger than –1; oppositely skewed baserates imply that the upper bound of the correlation is lower than +1.

[4] As the task was constructed to rule out prior expectancies, the sign of the correlation was conventionally defined to be "positive" when prebiotic diet is linked to many symptoms.

[5] In their original article, Meiser and Hewstone (2004) explained the observed stereotype in terms of an unwarranted inference from the two pairwise town-group and town-desirability correlations to the third correlation between group membership and desirability, which is just another way of framing the ecological correlation between group and desirability across contexts. Although rhetorically different, this earlier formulation is compatible with the present definition of PCs as baserate-driven inferences about correlations on the basis of information other than paired (x_i, y_i) observations.

[6] In contrast to the successive presentation format employed by Fiedler and Freytag (2004), the presentation of incomplete trivariate information by Meiser (2006) was not blocked into the presentation of town-group and town-desirability information, but sentences containing town-group and town-desirability information were displayed in random sequence.

[7] For example, given a zero contingency, the proportion of matching observations is .68 when both baserate distributions are .8 and .2 (i.e., $.8^2 + .2^2 = .64 +.04 = .68$), as compared with a matching proportion of .50 for equal baserates (i.e., $.5^2 + .5^2 = .25 +.25 = .50$).

[8] Statistically, the conditions of ecological bias are well understood in statistics (cf. Hammond, 1973). Ecological correlations diverge from individuating correlations when the independent variable impacts the ecological context, or when ecologies are dependent on the dependent variable.

Figure Captions

Figure 1: Illustrations of different relationships between a total correlation (big circle), ecological correlation (three dots), and partial correlations within ecologies (small circles).

Figure 2: Divergent ecological and individuating correlations between male versus female gender (M vs. F) and high versus low achievement (+ vs. –) across four school classes.

Figure 3: Even highly skewed marginals do not determine the sign of a contingency.

Figure 4: Panel a shows the trivariate stimulus distribution with negative partial correlation and positive total correlation of group membership and desirability. Panel b shows the true and estimated proportions of undesirable behaviors in a frequency-estimation task. Panel c shows the true and guessed proportions of behaviors stemming from Town X, from Group A within Town X (A|X) and from Group A within Town Y (A|Y) in a source-memory recognition task. Results from Meiser and Hewstone (2004, Experiment 1).

Figure 5: Trivariate stimulus distribution with negative partial correlation and zero total correlation of group membership and desirability. The full-information condition is displayed in Panel a, the incomplete-information condition is displayed in Panel b.

Figure 1

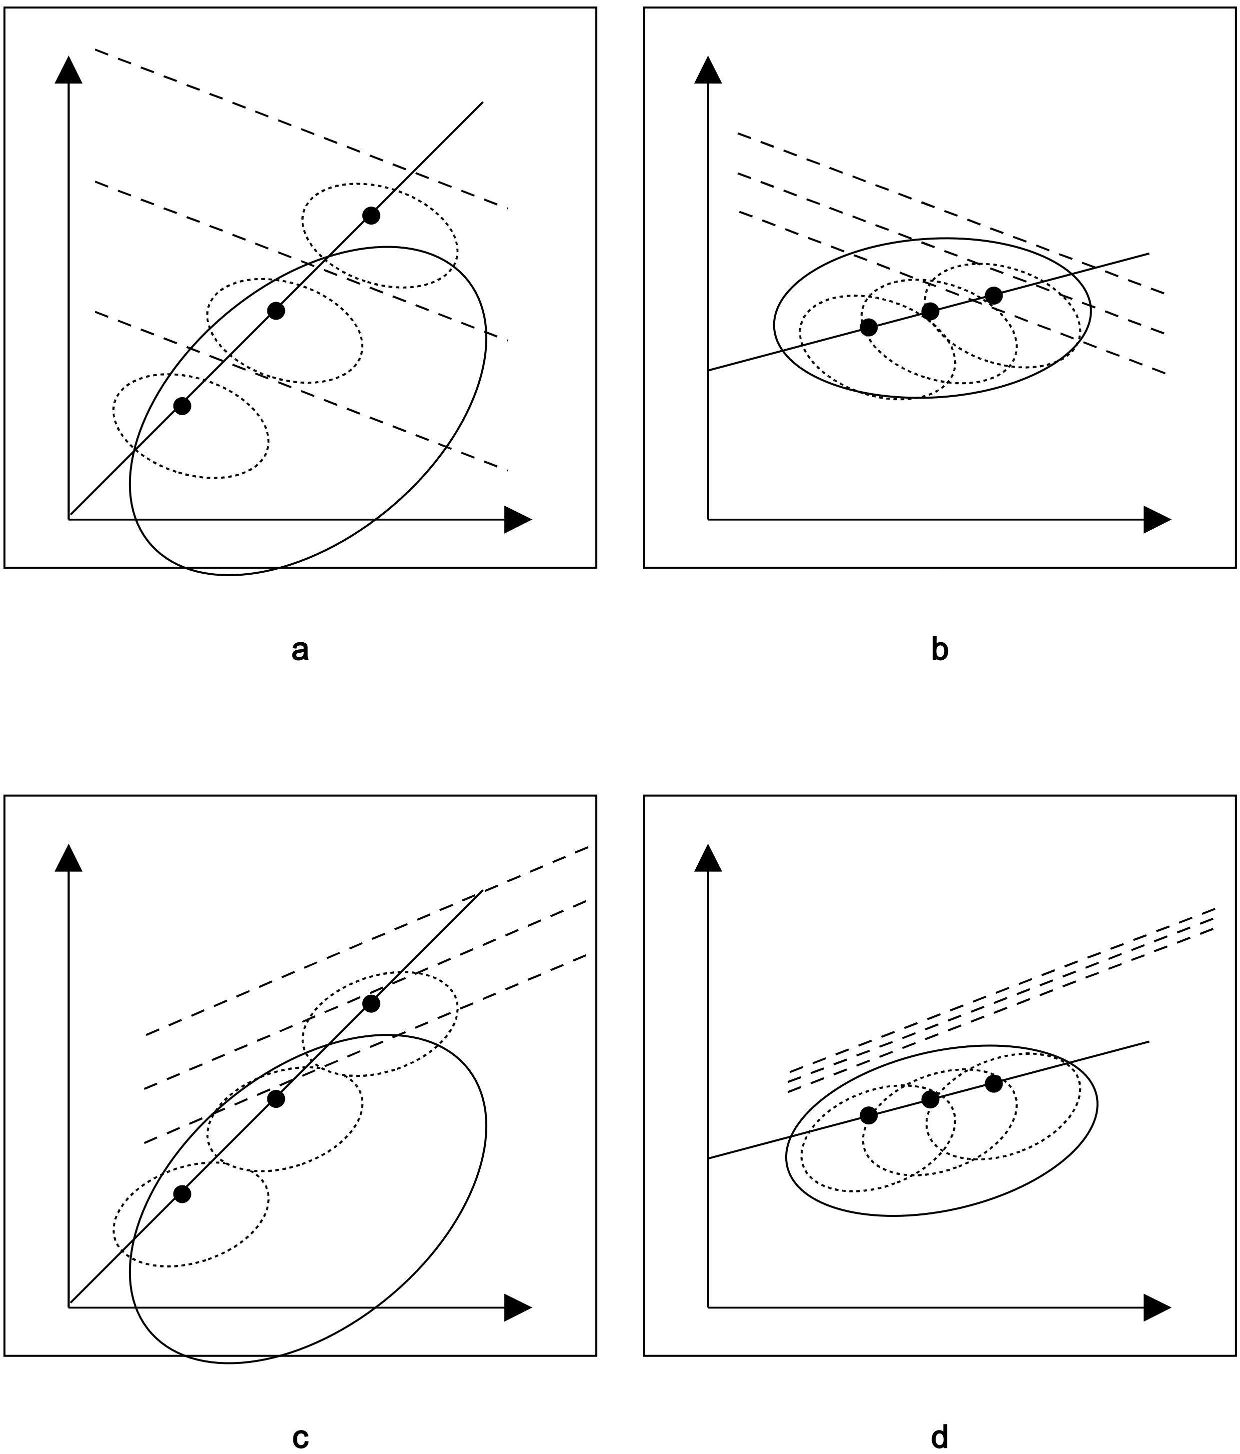

Figure 2

	+	-
M	12	6
F	6	0

	+	-
M	0	6
F	6	12

	+	-
M	12	6
F	6	0

	+	-
M	0	6
F	6	12

Figure 3

	Y+	Y-	
X+	?	?	36
X-	?	?	12
	36	12	

r = ?

a

	Y+	Y-	
X+	27	9	36
X-	9	3	12
	36	12	

r = 0

b

	Y+	Y-	
X+	33	3	[illegible]
X-	3	9	[illegible]
	[illegible]	[illegible]	

c

	Y+	Y-	
X+	24	12	36
X-	[illegible]	0	12
	[illegible]	[illegible]	

d

Figure 4

a) Trivariate Stimulus Distribution

	Town X		Town Y			
	Gr. A	Gr. B	Gr. A	Gr. B	Gr. A	Gr. B
Desirable	16	6	0	6	16	12
Undesirable	6	0	6	16	12	16

b) Proportions of Undesirable Behaviors

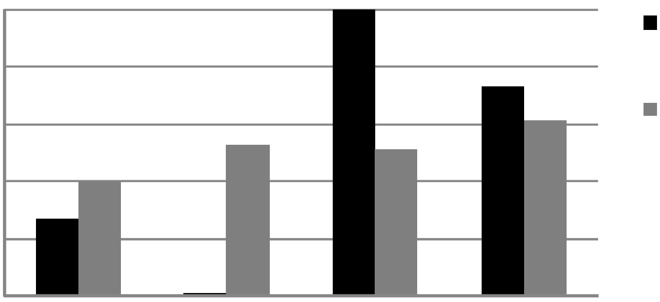

c) Proportions of Behaviors from Town X and Group A

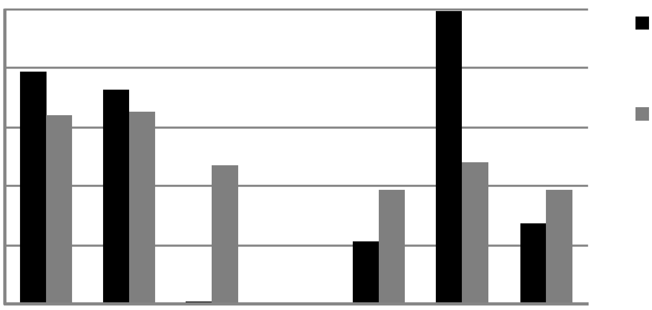

Figure 5

...rmation

	Town X		Town Y			
	Gr. A	Gr. B	Gr. A	Gr. B	Gr. A	Gr.
Desirable	16	8	0	8	16	16

...complete Trivariate Information

	Town X	Town Y
Group B	4	12

	Town X	Town Y
Undesirable	4	12

Nr.	Author	Title
08-40	Florian Kutzner Tobias Vogel Peter Freytag Klaus Fiedler	Pseudocontingencies in stereotype formation: extending illusory correlations
08-39	Wendelin Schnedler Adam Dominiak	Uncertainty Aversion and Preference for Randomization
08-38	Susanne Abele Sandra I. Vaughan-Parsons Garold Stasser	Information Flow and Influence during Collective Search, Discussion, and Choice
08-37	Peter Freytag Tobias Vogel Florian Kutzner Klaus Fiedler	The reproduction of base-rates promotes pseudocontingencies
08-36	Klaus Fiedler Peter Freytag Thorsten Meiser	Pseudocontingencies: An Integrative Account of an Intriguing Cognitive Illusion
08-35	Alexander Unger Dagmar Stahlberg	The Influence of Self-Regulation on Decision Making
08-34	Dirk Simons Nicole Zein	Kosten aus einer asymmetrischen Informationsverteilung zwischen Abschlussprüfer und Mandant
08-33	Michael Bremert Axel Schulten	The Impact of Supervisory Board Characteristics on Firm Performance
08-32	Jürgen Eichberger Ani Guerdjikova	Case-Based Expected Utility: Preferences over Actions and Data
08-31	Jürgen Eichberger Simon Grant Jean-Philippe Lefort	Neo-additive capacities and updating
08-30	Vladimir I. Danilov Gleb A. Koshevoy Christine Lang	Equilibria with indivisible goods and package-utilities

Nr.	Author	Title
08-29	Thomas Gschwend	Electoral System Change in Belgium 2003: Party Strategies and Voter Responses
08-28	Michael Herrmann	Expectations about Coalitions and Strategic Voting under Proportional Representation
08-27	Karl-Martin Ehrhart Marion Ott Susanne Abele	Auction Fever: Theory and Experimental Evidence
08-26	Michel Regenwetter	Perspectives on Preference Aggregation
08-25	Daniel M. Bernstein Michael E. Rudd Edgar Erdfelder Ryan Godfrey Elizabeth F. Loftus	The Revelation Effect for Autobiographical Memory: A Mixture-Model Analysis
08-24	Markus Glaser Florencio Lopez de Silanes Zacharias Sautner	Looking Inside a Conglomerate: Efficiency of Internal Capital Allocation and Managerial Power Within a Firm
08-23	Markus Glaser Martin Weber	Financial Literacy und Anlegerverhalten
08-22	Markus Glaser Thomas Langer Jens Reynders Martin Weber	Scale Dependence of Overconfidence in Stock Market Volatility Forecasts
08-21	Patrick A. Müller Rainer Greifeneder Dagmar Stahlberg Kees Van den Bos Herbert Bless	The Role of Procedural Fairness in Trust and Trustful Behavior
08-20	Patrick A. Müller Rainer Greifeneder Dagmar Stahlberg Kees Van den Bos Herbert Bless	The Influence of Accessibility Experiences on the Willingness to Cooperate in Negotiations

Nr.	Author	Title
08-19	Jana Janßen Patrick A. Müller Rainer Greifeneder	A field study on the role of ease-of-retrieval in procedural justice judgments
08-18	Rainer Greifeneder Benjamin Scheibehenne	When choosing is difficult: Complexity and choice-overload
08-17	Paul Grout Wendelin Schnedler	Non-Profit Organizations in a Bureaucratic Environment
08-16	Clemens Kroneberg Isolde Heintze Guido Mehlkop	On shoplifting and tax fraud: An action-theoretic analysis of crime
08-15	Hermann Jahnke Dirk Simons	A rationale for the payback criterion
08-14	Peter Dürsch Jörg Oechssler Radovan Vadovic	Sick Pay Provision in Experimental Labor Markets
08-13	Carsten Schmidt Martin Strobel Henning Oskar Volkland	Accuracy, Certainty and Surprise - A Prediction Market on the Outcome of the 2002 FIFA World Cup
08-12	Mathias Sommer	Understanding the trends in income, consumption and
08-11	Hans Gersbach Hans Haller	Club Theory and Household Formation
08-10	Michael F. Meffert Thomas Gschwend	Strategic Voting in Multiparty Systems: A Group Experiment
08-09	Jens Wüstemann Jannis Bischof	Ausweis von Finanzinstrumenten in europäischen Bankbilanzen nach IFRS: Normative Erkenntnisse empirischer Befunde
08-08	Jürgen Eichberger David Kelsey	Are the Treasures of Game Theory Ambiguous?
08-07	Jürgen Eichberger Ani Guerdjikova	Multiple Priors as Similarity Weighted Frequencies