SonderForschungsBereich 504

Rationalitätskonzepte,
Entscheidungsverhalten und
ökonomische Modellierung

No. 08-40

Pseudocontingencies in stereotype formation: extending illusory correlations

Florian Kutzner*
and Tobias Vogel**
and Peter Freytag***
and Klaus Fiedler****

January 2009

Correspondence concerning this paper should be addressed to Florian Kutzner, Department of Psychology, University of Heidelberg, Hauptstrasse 47-51, 69117 Heidelberg, Germany, Phone +49-6221-547366, Fax +49-6221-547745, Email: florian.kutzner@psychologie.uni-heidelberg.de

*Sonderforschungsbereich 504, email: florian.kutzner@psychologie.uni-heidelberg.de

**Sonderforschungsbereich 504, email:

***Universität Heidelberg, email:

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

Universität Mannheim
L 13,15
68131 Mannheim

Pseudocontingencies in stereotype formation:

extending illusory correlations

Florian Kutzner, Tobias Vogel, Peter Freytag & Klaus Fiedler

Universität Heidelberg

Running head: PC's in stereotype formation.

Correspondence concerning this paper should be addressed to Florian Kutzner, Department of Psychology, University of Heidelberg, Hauptstrasse 47-51, 69117 Heidelberg, Germany, Phone +49-6221-547366, Fax +49-6221-547745, Email: florian.kutzner@psychologie.uni-heidelberg.de.

Abstract

Under the notion of illusory correlations, simple learning paradigms (e.g. Hamilton & Gifford, 1976) have been used to study the formation of stereotypes that discriminate between majorities and minorities. In the present paper, limitations of this approach in terms of theoretical explanations and empirical evidence are addressed. Theoretically, we propose pseudocontingencies (PCs, Fiedler, Freytag & Meiser, 2008) as a more robust mechanism behind illusory correlations. In contrast to previous explanations, PCs can explain illusory correlations when groups are never paired with valence. Empirically, we replicate earlier findings, i.e. that the more frequently observed group, the majority, is evaluated more in line with the more frequently observed valence. Crucially, we extend the empirical evidence in that illusory correlations prove robust over a very large number of observations (320) and under increasingly interactive task conditions, involving predictions of valence (Experiment 2) and reinforcement-learning conditions (Experiment 3). The latter provided evidence for illusory correlations on a new measure, participants' predictions. These predictions reflect the expectations about the valence associated with majority and minority and might well affect real life behavior. The discussion focuses on possible reasons for why PCs are used in stereotypic judgments.

191 words

Keywords: pseudocontingency, skewed base rates, illusory correlation, probabilistic contingency learning, reinforcement learning, stereotype formation

Pseudocontigencies in frequency based illusory correlations

Ever since the seminal work of Hamilton and Gifford (1976) on frequency-based illusory correlations (ICs), social psychologists have been intrigued by the notion that group stereotypes may simply depend on the distribution of information in the social environment, independent of any sentiment or real group conflict. In their experimental analog of stereotype formation, participants first observed a series of statements describing positive or negative behaviors shown by members of two fictitious groups. The series of stimulus behaviors had three key characteristics (e.g. Hamilton & Gifford, 1976, Exp.1): (a) positive behaviors appeared more often than negative ones, in 27 and 12 out of 39 instances, respectively, (b) one group (i.e., the majority) appeared more often than the other, in 26 and 13 instances, and (c) the ratio of positive to negative behaviors was the same (viz., 2:1) for both groups, setting the correlation between groups and valence to zero. Therefore, there was no apparent reason to evaluate the two groups differentially. Nevertheless, subsequent group impression ratings, frequency estimates and recall responses reflected more positive evaluations of the majority than the minority.

This phenomenon of illusory correlation was consistently replicated in numerous studies (for a review see Mullen & Johnson, 1990). If positive behaviors prevailed in the stimulus series, the majority appeared more positive than the minority. In contrast, if negative behaviors prevailed, the majority appeared more negative. In general, an illusory correlation was formed between the more frequently presented group and the more prevalent valence.

To account for this clear-cut evidence on the illusory-correlation phenomenon, different theoretical explanations have been offered (Fiedler, 2000). Originally, Hamilton and Gifford (1976) had assumed that the most infrequent group-behavior combination (e.g., negative minority behavior in the original design with mostly positive behavior) is most distinctive and therefore enjoys a memory advantage over the other combinations. This could account for the relatively negative evaluation of the minority. This explanation was not supported in later studies using more refined memory analyses (Fiedler, Russer & Gramm, 1993; Klauer & Meiser, 2000). These studies instead suggested an explanation in terms of unequal learning opportunities, pointing out that (at least) twice as many learning trials are available for the majority as for the minority (Smith, 1991; Fiedler, 1996).

Theoretical and Empirical Limitations

However, existing explanations and evidence on illusory correlations converge in a number of theoretical and empirical limitations. First, and most crucially, it is taken for granted in nearly all articles published so far that the underlying cognitive process involves correlation assessment proper. It is presupposed that participants form a stereotypical correlation from *joint observations* of group members and positive or negative behaviors. Thus, for any positive or negative behavior encountered, it is essential to know the group reference. If this condition is not met, no correlation can be determined. In our view, explaining illusory correlations with correlation assessment proper unnecessarily restricts the range of real life situations that can be explained.

Second, all pertinent research so far has operationalized the formation of illusory stereotypes as a completely passive, *experimenter-determined observation process.* Participants are not allowed to make active inferences or predictions and they do not receive dynamic feedback to their own inferences and intermediate judgments. This restriction is certainly not representative of the way in which stereotypes are learned in reality. Assuming a more active or interactive process, biased inferences probably have unwanted consequences that might “correct” the illusion. For example, an inflated rate of predicting negative behaviors for the minority as compared to majority should be disconfirmed by corrective feedback, mostly portraying the minority as not that negative. If it can be shown that the illusion will nevertheless appear in such a dynamic feedback environment, this would provide particularly strong evidence for its external validity.

A third limitation refers to the long-term stability of stereotypes. Virtually all published experiments include a limited number of observations between 36 and 48 behavior descriptions (for an exception see Fiedler, Hemmeter, Hofmann, 1974). It is unclear, both theoretically and empirically (cf. Mullen & Johnson, 1990, p. 24), whether the illusion would survive a more *extended learning process*. After all, both the distinctiveness of rare events, and the lesser opportunity to learn about the minority, might loose their impact when the number of observations increases substantially. However, if the illusory correlation phenomenon is to account for enduring stereotypes, it is essential to demonstrate that illusory correlations remain stable across a large number of learning trials.

Goals of Present Research

In the present research, we try to tackle all three limitations. We demonstrate that a cognitive process more widely applicable than correlation assessment, denoted pseudocontingencies (PC), can produce illusory discrimination between majority and minority groups. Moreover, this process is shown to result in a stable illusion over a very large number of observations. And finally, this holds even for interactive task conditions involving predictions and corrective feedback. First, let us introduce the notion of pseudocontingencies (PCs) as an alternative process to correlation judgments.

Pseudocontingencies in illusory correlations

PCs are correlation inferences driven by base rate information. In the context illusory correlations, base rates refer to the occurrence rates of the two groups and the prevalence of positive and negative behaviors. In this context, a PC is implicated when one group is more prevalent than the other, and when one valence is more prevalent than the other. The alignment of the group and valence base rates determines the PC inference: *observations of the frequent (rare) value on one attribute coincide with the frequent (rare) value on the other attribute*.

To illustrate the differences between PCs and correlation assessment proper, consider the stimulus distribution depicted in Table 1. Over as many as 320 observations, the joint frequencies of positive and negative behaviors in the purple, majority group are 180 (+) and 60 (–), respectively. The corresponding joint frequencies for the orange, minority group are 60 (+) and 20 (–), respectively. A normatively correct correlation inference would have to use these cell frequencies (Allan, 1993; Fiedler, 2000), which indicate the distribution of one variable as a function of different levels of the other variable (e.g., the rate of positive behaviors conditional on a group). In contrast, PCs take the alignment of base rates for a correlation. The (statistically unwarranted) tendency to infer that valence (+ vs. –) is contingent on groups (purple vs. orange) is derived from an alignment of the skewed base-rate distributions in the marginals of the frequency table. Given unequal base rates of 240 and 80 for groups and 240 and 80 for behaviors, respectively, the more frequent group is aligned with the more frequent valence level, thus resulting in the PC inference that the purple group is associated with + and, likewise, the orange group is associated with –.

A PC account of more favorable majority than minority judgments applies to a wider range of real life learning conditions than a traditional illusory-correlation account. PC inferences are still possible when the number of observations is too small to allow for correlation assessment proper. Fewer observations are needed to detect unequal group and valence base rates than to detect a correlation, that is, the differential rates of positive and negative valence as a function of groups.

Another reason why PCs can be inferred when correlation assessment is impossible, in addition to limited number of observations, is that information about groups and valence may be dissociated. Very often, the social environment does not provide joint observations about group membership and valence. We may rather learn that a majority group is larger than a minority group on one occasion and, on a completely different occasion, that positive behavior is more prevalent than negative behavior. They may not receive joint information about whether positive and negative behavior is paired with the majority or minority group. In this case, the correlation is unspecified, although PCs are still possible. Indeed, there is evidence on illusory stereotypes (McGarty, Haslam, Turner & Oakes, 1993) and on contingency illusions more generally (Fiedler & Freytag, 2004), showing unjustified discrimination between groups in the absence of joint observations. Thus, PCs explain illusory discrimination between majorities and minorities under a wider range of environmental conditions than traditional accounts of illusory correlations.

Present experiments and modifications

All three experiments reported in this article share the same basic set up. They consist of a learning phase and, immediately afterwards, an evaluation phase. On every trial of the learning phase, the label for one of two groups is presented, followed by valenced behavior descriptions. The joint frequencies of group labels and behaviors appear in Table 1. We conserve all crucial properties of stimulus distributions used to demonstrate illusory correlations, namely, skewed base rates of group and valence and a zero correlation between the attributes. In the subsequent evaluation phase, attitude ratings towards the groups are assessed independently for the more and less frequently observed group.

From the first to the last experiment, we gradually introduce three modifications of the standard frequency-based illusory correlation paradigm. In doing so, we address the three limitations outlined above: whether illusory correlations (a) show long-term stability, (b) are maintained given the participants' active involvement and their corrective feedback, and (c) persist when only PCs provide a reasonable explanation.

Extended stimulus series. In all three experiments, we increase the number of stimulus observations far beyond the amount of information used in prior research. Participants experience 320 trials in which unidentified members of either the majority (240 trials) or the minority (80 trials) show a positive (240) or a negative (80) behavior.

The purpose for extending the stimulus series is, obviously, to demonstrate the long-term stability of the illusion. However, if the illusory discrimination of groups turns out to persist over more than 300 trials, this would not only provide evidence for temporal stability. It would also highlight the fact that the illusion survives in the face of a very large sample of observations that converges to a plain zero contingency. Because PCs reflect only aligned base rates, even a greatly extended stimulus series that maintains the base rates should not reduce the illusory correlation effect. Thus, extending the domain of illusory correlations to the case of very long and exhaustive intergroup experience testifies to the robustness of the phenomenon and at the same time supports a PC-based interpretation.

Predictions of valence. Whereas the extension of the stimulus series is the novel feature of Experiment 1, Experiments 2 and 3 replace the passive observation task that characterizes previous studies by a more active or interactive prediction task. From the first trial on, participants are asked to predict the valence of the upcoming behavior given the group membership as predictor.

This change in the task setting is meant to reflect how stereotypes are usually formed and used in real life. We refer to situations in which the group membership of a person that is known beforehand is used to predict the valence of that person's behavior. In Experiment 2, we demonstrate that the illusory correlations generalize from purely memory-based judgments analyzed in virtually all previous studies to a new task setting involving the active formation of expectancies and continuous predictions. Group discrimination in such an active task setting would disconfirm the suspicion that illusory correlations are confined to memory instructions as opposed to instructions encouraging impression formation during stimulus encoding (REF xxxx).

Corrective feedback. In Experiment 3 participants additionally receive corrective feedback for their predictions. That is, for every correct prediction, participants receive a pay-off of 3 cents; for every incorrect prediction, they loose the same amount of 3 cents. The purpose of introducing payoffs was twofold. First, we model situations in which stereotype formation and usage is not only interactive, but also of personal relevance. Second, for the reasons discussed below, this modification might yield even stronger evidence for the robustness of illusory correlations and for an interpretation in terms of PCs.

Providing feedback and reinforcements participants' predictions creates operant learning conditions. Under these conditions, the predictions themselves provide a natural online measure of correlation judgments (Allen, 1993). An illusory correlation effect is reflected in a higher rate of predicting the frequent valence for the frequent group than for the infrequent group. In an operant task context, such an illusory prediction bias reduces the overall amount of reinforcement. The maximal pay-off would be obtained by always predicting the more frequent valence, irrespective of the group. If the participants' online predictions nevertheless discriminate between groups, this would reflect a particularly robust manifestation of an illusory correlation effect, which persists despite a counter-active pay-off scheme.

Moreover, the online manifestation of the illusion can be determined at any segment of the operant-learning process. Monitoring the development of the illusion across the entire stimulus series allows for extrapolating inferences. If the size of the illusion reaches a stable asymptote, this would provide particularly strong evidence for robustness and stability. In a non-social matching-to-sample paradigm, we already obtained first evidence for a similar illusion after extended corrective feedback (Kutzner, Freytag, Vogel, & Fiedler, 2008).

To summarize, the present research extends illusory correlations, both theoretically and empirically. We try to provide evidence for PCs as a theoretical explanation. Within such an explanatory framework, illusory correlations can be predicted even under conditions of limited information that does not warrant genuine correlation assessment. Additionally, we try to show the robustness of the illusion my introducing various ecologically plausible modifications of the standard illusory correlation paradigm. We greatly increase the number of stimulus observations and involve participants in active predictions and by providing them incentives for accurate responses.

Hypotheses

We hypothesize that even after many observations (i.e. 320) the more frequently observed group (i.e. the majority) will be evaluated more in line with the more frequent valence of behaviors than the minority. When participants make predictions under operant-learning conditions, we expect a tendency to more frequently predict the frequent type of behavior for the majority as compared to the minority (Exp. 3). This effect should be manifested in the predictions, notwithstanding the monetary incentives for accurate predictions. As we expect the illusions in the online predictions and in the final judgments and evaluations to reflect the same basic process, we expect the novel bias in the online predictions to correlate systematically with the memory-based judgments and evaluations at the end of all observations.

******

Experiment 1

In Experiment 1 we closely replicated standard frequency based illusory correlation paradigms. However, we extended the number of observations to 320.

Method

Participants and Design. Twenty-one undergraduate students (11 female, 10 male) from the University of Mannheim participated in the study for a compensation of 4€. Subjects were randomly assigned to one of the two conditions with different types of frequent behaviors. Either positive or negative behaviors were more frequent, at a ratio of three to one. Every subject was exposed to the two different types of groups, the majority and the minority. Observations relating to the majority were also three times as frequent as those relating to the minority (cf. Table 1). The groups did not differ in the ratio of good and bad behaviors. This resulted in a 2 (frequent behavior positive vs. negative) x 2 (majority vs. minority) mixed design with repeated measures on the latter factor.

Procedure. Participants were informed that they would have to form an accurate impression about two groups over a series of observations. To move from trial to trial participants had to press a key on their keyboard. We referred to the groups as the “Purple group” and the “Orange group.” When the labels appeared on the screen they were accompanied by a corresponding hue as background of the screen and the keyboard was locked for 1.0 s. When subjects had pressed one of the keys, valenced behavioral descriptions appeared next to the group label for 3.0 s, creating joint observations of group label and valence. We used 128 behavioral descriptions that had been pretested for their valence. For example, a negatively valenced description would read “A member the purple/orange group uses vast amounts of pesticides in his garden,” a positive “A member the orange/orange group can laugh about his own weaknesses.” After 0.3 s without group label or behavioral descriptions visible on the screen, the next group label appeared.

To compose the sequence of 320 trials, for every subject the computer generated 20 random sequences each consisting of 16 trials that retained the statistical properties of the distribution in Table 1. These 20 sequences were adjoined to produce the full 320 trial sequence.

After this training phase, participants indicated their overall impression of the two groups rating both groups on likeability and their intentions to spend time with members of each group. Computers delivered the instructions, controlled stimulus presentation, and recorded participants’ evaluations.

Results and Discussion

We submitted the evaluations of the groups to a two-factor repeated-measures analysis of variance with type of frequent behavior as the between-subject factor (positive vs. negative) and type of group as the within-subject factor (majority vs. minority). The analysis revealed a marginally significant main effect for the valence factor ($F(1, 19) = 3.59$, p=.07, η^2 = .16) and a marginal interaction ($F(1, 19) = 3.13$, p=.09, η^2 = .14). As Figure 1 shows, the more frequent group, the majority, was preferred if the more frequent behavior was positive and the less frequent group, the minority, was preferred when the less frequent behavior was positive. Looking at the average effect size of r^2 = .14, the results closely match those of Mullen and Johnson's (1990) meta-analysis (r^2 = .12).

Insert Figure 1 about here

As compared to standard illusory correlation paradigms, in Experiment 1 we increased the number of observations to address the long-term stability of the illusion. We found that still after 320 trials a preference persisted for the majority when positive behavior was more frequent and for the minority when negative behavior was more frequent. These results are first evidence for the long-term stability of the illusory correlation and with that for PCs as a possible explanation.

Experiment 2

Experiment 2 was a first step to move away from the completely passive, experimenter-determined observation task of merely observing group members and behaviors. Before the behaviors could be observed, participants were asked to predict their valence given the group membership as predictor. To mimic that stereotype judgments might be made based on very few observations, we asked participants to make predictions from the first trial on.

Method

Participants and design. Twenty-three undergraduate students (15 female, 8 male) from the University of Mannheim participated in the study for a compensation of 4€. We employed the same stimulus distribution (cf. Table 1) as in Experiment 1 with the same 2 (frequent behavior positive vs. negative) x 2 (majority vs. minority) mixed design with repeated measures on the latter factor.

Procedure. As in Experiment 1, participants were informed that they would have to form an accurate impression about two groups over a series of observations. This time however, they were additionally instructed that to do so they had to predict whether group members would show positive or negative behaviors on the next observation. They did so by pressing one of two response keys on the left or the right side of a German computer keyboard. The meaning of the keys remained constant over the course of the experiment but was counterbalanced between subjects. After the prediction, the predefined behavioral description appeared on the screen together with the group label. In contrast to the next study, subjects' predictions were not followed by explicit feedback as to whether their predictions were correct or incorrect and had no consequences for their reward. The same group labels and behavioral descriptions as in Experiment 1 were used.

After this training phase, participants indicated their overall impression of the two groups by answering the same evaluation questions as in Experiment 1. Again, computers delivered the instructions, controlled the stimulus presentation of the 320 observations, and recorded participants' choices.

Results and Discussion

We analyzed the predictions by forming blocks of 32 trials. We computed the proportion of predictions for the more frequently presented valence conditional on the type of group presented before. These proportions were submitted to a three-factor repeated-measures analysis of variance with valence of frequent behavior as the between-subject factor (positive vs. negative) and block (1-10) and type of group as within-subject factors (majority vs. minority). The analysis revealed a strong main effect for the between subject factor, type of frequent behavior ($F(1, 21)=11.85$, $p<.001$, $\eta^2 = .88$) indicating a more pronounced tendency to predict the frequent type behavior if the frequent type was positive (all interactions with type of frequent behavior F's < 1.5, ns.). No other effect was significant, neither the main effect for block, $F(9, 13) = 1.97$, $p=.128$, $\eta^2 = .58$, nor the main effect for type of group, F(1, 21)< 1, nor the interaction between type of group and block ($F(9,13) < 1$). Figure 2 shows the mean choice proportions for the more frequent behavior conditional on the type of group.

Insert Figure 2 about here

We submitted the evaluations of the groups obtained after predictions to a two-factor repeated-measures analysis of variance with type of frequent behavior as the between-subject factor (positive vs. negative) and type of group as the within-subject factor (majority vs. minority). The analysis revealed a main effect for the valence factor ($F(1, 21) = 6.72$, $p<.05$, $\eta^2 = .24$) and an interaction ($F(1, 21) = 5.49$, $p<.05$, $\eta^2 = .21$). As Figure 3 shows, the more frequent group, the majority, was preferred if the more frequent behavior was positive and the less frequent group, the minority, was preferred when the less frequent behavior was positive. Looking at the average effect size $r^2 = .21$, this time the illusory correlations are even stronger than in Mullen and Johnson's (1990) meta-analysis ($r^2 = .12$).

Insert Figure 3 about here

To predict the illusory correlation obtained from the evaluations by the illusion inherent to the predictions we computed difference scores. For the evaluations we obtained an illusory correlation measure by subtracting the evaluation of the minority from that of the majority if the frequent valence of behaviors was positive and vice versa if the frequent valence was negative. To obtain an illusory correlation measure for the predictions, we computed the difference scores for each of the 10 blocks. We subtracted the proportions of predictions for the frequent type of behavior after the majority from the proportions after the minority. We built an average prediction correlation judgment by averaging the 10 subsequent difference scores (α=.94). A correlation analysis revealed that the prediction correlation judgment did not correlate with the illusory discrimination in evaluations of the groups ($r(23)$=.16, ns.).

In Experiment 2, participants were actively involved in the task by predicting the valence of behaviors, again, over an extended number of observations. The illusory correlations also proved robust against this ecologically plausible modification. Participants exhibited a preference for majority when the frequent behavior was more frequent and for the minority when negative behavior was more frequent. These results extend the evidence that PCs might underlie illusory correlations under interactive task conditions.

Experiment 3

In the last experiment, we built upon Experiment 2 and further increased participants' involvement by reinforcing correct predictions. Crucially, this enabled us to monitor the acquisitions process of the illusory correlation and to address the question whether further observations would have reduced the illusion.

Method

Participants and design. Thirty-seven undergraduate students (21 female, 16 male) from the Universities of Heidelberg and Mannheim participated in the study for performance-contingent reward that averaged around 4€. Design and stimulus distributions were identical to Experiments 1 and 2. Subjects were randomly assigned to the 2 (frequent behavior positive vs. negative) x 2 (majority vs. minority) mixed design with repeated measures on the latter factor.

Procedure. This time, participants were informed that they could win money making correct predictions about the valence of the member's behavior. If the predefined valence was predicted, 0.03€ were added to the subjects account. In the case predictions did not match the valence of the behavioral description, 0.03€ were subtracted. The updated account value and the last change were constantly displayed in the lower right corner of the screen.

After this training phase, participants again indicated their overall impression of the two groups. The rest of the procedure was identical to the previous experiments.

Results and Discussion

We analyzed the predictions by forming blocks of 32 trials. We computed the proportion of predictions for the more frequently presented valence conditional on the type of group presented before. These proportions were submitted to a three-factor repeated-measures analysis of variance with valence of frequent behavior as the between-subject factor (positive vs. negative) and block (1-10) and type of group as within-subject factors (majority vs. minority). The analysis revealed a marginal main effect for the between subject factor, type of frequent behavior ($F(1,35)=3.88$, $p=.06$, $\eta^2 = .10$) indicating a more pronounced tendency to predict the frequent type behavior if the frequent type was positive (all interactions with type of frequent behavior *F*'s < 1.4, ns.) and a strong main effect for block, $F(9, 27) = 15.70$, $p<.001$, $\eta^2 = .84$. Crucially, we obtained a main effect for type of group, $F(1, 35) = 14.45$, $p<.01$, $\eta^2 = .21$, that did not depend on the block, i.e. there was no interaction between type of group and block ($F < 1$). Figure 4 shows the mean choice proportions for the more frequent behavior conditional on the type of group.

Additionally, we performed two analyses to address a possible change in the size of the type-of-group effect. First, we analyzed the difference in predictions rates in the last block. A repeated measures ANOVA revealed that still for the last 32 predictions there was the tendency to predict more of the frequent behavior for the majority than for the minority ($F(1,36)=5.17$, $p<.05$, $\eta^2 = .13$). Second, we compared the difference in predictions rates for the first and last block of 32 trials. With just these two blocks, a 2-factor repeated measures ANOVA showed a strong main effect for block, $F(1, 36) = 102.70$, $p<.001$, $\eta^2 = .74$, a main effect for type of group, $F(1, 36) = 10.62$, $p<.01$, $\eta^2 = .23$, and, crucially, no interaction between type of group and block ($F < .12$).

Insert Figure 4 about here

As before, we also submitted the evaluations of the groups to a two-factor repeated-measures analysis of variance with type of frequent behavior as the between-subject factor (positive vs. negative) and type of group as the within-subject factor (majority vs. minority). The analysis revealed a main effect for the valence factor ($F(1, 35) = 7.64$, $p<.01$, $\eta^2 = .18$) and an interaction ($F(1, 35) = 19.64$, $p<.001$, $\eta^2 = .36$). As Figure 5 shows, the more frequent group, the majority, was preferred if the more frequent behavior was positive and the less frequent group, the minority, was preferred when the less frequent behavior was positive.

Insert Figure 5 about here

To predict the illusory correlation obtained from the evaluations by the illusion inherent to the predictions we computed difference scores. For the evaluations we obtained an illusory correlation measure by subtracting the evaluation of the minority from that of the majority if the frequent valence of behaviors was positive and vice versa if the frequent valence was negative. To obtain an illusory correlation measure for the predictions two steps were necessary. First, we computed the difference scores for each of the 10 blocks. We subtracted the proportions of predictions for the frequent type of behavior after the majority from the proportions after the minority. Then, we analyzed the internal consistency of these 10 subsequent correlation judgments. Corrected item-total correlations revealed that the first two judgments did only correlate weakly with the total (r's$<.25$). Therefore, we constructed two average correlation judgments, for the first two blocks ($r(37)=.68$, $p<.001$) and the following eight blocks ($\alpha=.85$). The scales did not correlate ($r(37)=.08$, ns.).

The main regression analysis revealed that both correlation judgments together explained 26% of the variance in the illusory discrimination in evaluations of the groups ($F(2,34)=5.99$, $p<.01$). Both, early and late correlation judgments separately predicted the illusory discrimination ($\beta=.33$, $p<.05$ and $\beta=.35$, $p<.05$, for the first two blocks and the following eight respectively).

In this last experiment, we increased the personal relevance of correct predictions by introducing reinforcement-learning conditions. This provides a completely new measure of illusory correlations, participants' predictions of valence given the groups. Parallel to the evaluations of the groups, this behavioral measure also revealed stable illusory correlations. Interestingly, though unpredicted, early and late correlation judgments in the predictions did not correlate. This suggests that participants differ in when they most show the illusion.

Additionally, given the characteristics of the acquisitions process it is unlikely that further observations would have reduced the illusion. Most importantly, because the size of the illusion did not depend on when it was assessed, after 32 or 320 observations. Also supporting the idea that learning had stabilized, from Figure 3 it seems that predictions rates did remain constant after half of the observations. Thus, as predicted by the PC, illusory correlations show long-term stability even in the face of personally relevant corrective feedback.

General Discussion

In the present research we extend the validity of illusory correlations (Hamilton & Gifford, 1976) as an analogue of how majority-minority stereotypes are formed in real life. On the one hand, we provide evidence for an explanation, namely pseudocontingencies (PCs, Fiedler, Freytag & Meiser, 2008), that captures a wider range of ecological conditions under which these stereotypes might be formed. On the other hand, we generate additional evidence for illusory correlations under ecologically plausible modifications of the standard task.

PCs in illusory correlations

PCs are correlation judgments that rely on aligned base rates, in other words on the fact that for the two attributes to be correlated some observations are more frequent than others. In illusory correlations, the frequent attributes are (a) the majority and (b) either positive or negative valence of behavior. Given these aligned base rates, PCs link the frequent observations and the infrequent observations, thereby creating the illusion of a correlation. Unlike explanations that rely on correlation assessment proper, PCs predict this illusory correlation disregarding information about the pairings of group members and valence.

Only PCs predict illusory correlations in situations where there is only limited information about the majority and, even more plausibly, about the minority. For example, if no minority members have been observed performing clearly positive or negative behaviors, no correlation between group and valence can be determined. Beyond plausibility, there is suggestive evidence that people actually make stereotypic judgments without joint observations ((McGarty et al., 1993; Fiedler & Freytag, 2004). But even if joint observations are available, but information is limited in terms of the number of observations, PCs enjoy an advantage. Every observation entails all the information necessary to update both base rate estimates, of groups and valence, but only half the information needed to estimate the two conditional base rates of valence given one or the other group. Thus, PCs explain majority-minority stereotype formation even if very little is known about the groups in question.

What might seem paradoxical given this limited-information advantage, PCs also imply that even extensive experience with group members paired with valence will not eliminate the illusion. This is due to the fact that PCs are only sensitive to base rate information, i.e. how often groups and positive or negative behaviors are observed. In this clear-cut form, only PCs make this prediction. For example, Mullen and Johnson (1990), favoring distinctiveness-based explanation, concluded that either stronger or weaker illusory correlations might have occurred with an increasing number of observations (p.24).

Consequently, we conducted experiments in which we drastically increased the number of observations. All three experiments demonstrated illusory correlations in stereotype formation regarding two fictitious groups even after 320 observations. That is we found more positive attitude ratings for the majority if the "majority" of behaviors was positive and for the minority of the "minority" of behaviors was positive. Crucially, based on the monitoring of the learning process of the stereotype in Experiment 3, we argue that further increasing the number of observations would not have eliminated the illusion. This provides strong evidence that aligned base rates, i.e. PCs, not joint observations of groups and valence explain part of the illusory correlation.

References

Allan, L. G. (1993). Human contingency judgments: Rule based or associative? *Psychological Bulletin, 114(3)*, 435-448.

Denrell, J. (2005). Why Most People Disapprove of Me: Experience Sampling in Impression Formation. *Psychological Review, 112(4)*, 951-978.

Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review, 103(1)*, 193-214.

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review, 107*(4), 659-676.

Fiedler, K., & Freytag, P. (2004) Pseudocontingencies. *Journal of Personality and Social Psychology, 87(4)*, 453-467.

Fiedler, K., Russer, S., & Gramm, K. (1993). Illusory correlations and memory performance. *Journal of Experimental Social Psychology, 29(2)*, 111-136.

Fiedler, K., Freytag, P., & Meiser, T. (2008). *Psychological Review.*

Hamilton, D. L., & Gifford, R. K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12(4)*, 392-407.

Klauer, K. C., & Meiser, T. (2000). A source-monitoring analysis of illusory correlations. *Personality and Social Psychology Bulletin, 26(9)*, 1074-1093.

Kutzner, F., Freytag, P., Vogel, T., & Fiedler, K. (2008). Base-rate neglect as a function of base rates in probabilistic contingency learning. *Journal of the Experimental Analysis of Behavior, 90(1)*, 23-32.

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(4)*, 391-410.

Mullen, B., & Johnson, C. (1990). Distinctiveness-based illusory correlations and stereotyping: A meta-analytic integration. *British Journal of Social Psychology, 29(1)*, 11-27.

Smith, E. R. (1991). Illusory correlation in a simulated exemplar-based memory. *Journal of Experimental Social Psychology, 27(2)*, 107-123.

Table 1

Table 1. Contingency tables indicating the predetermined distributions used in the experiments. The base rates of both predictors (group labels) and correct predictions (valence of behavior) were skewed at a ratio of 3:1.

Experiments 1 & 2

		Correct prediction		
		Positive behavior	Negative behavior	
Predictor	Purple group	180	60	240
	Orange group	60	20	80
		240	80	320

Figure 1

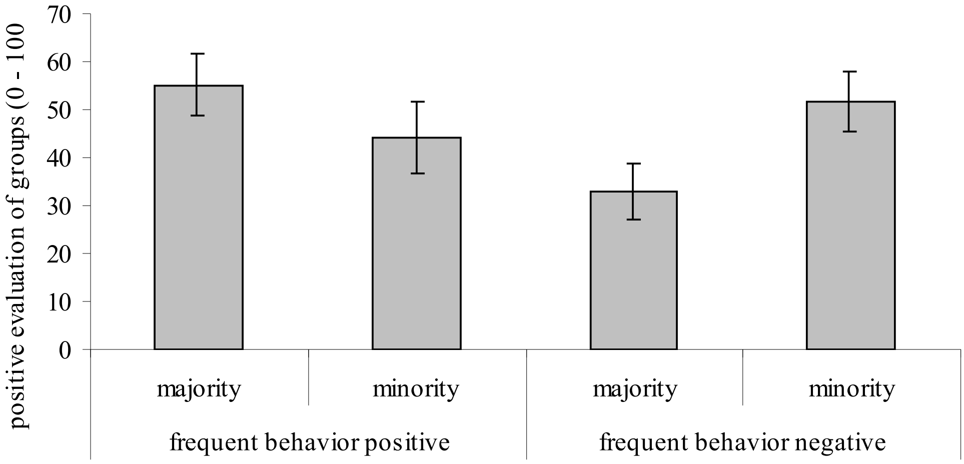

Figure 1. Mean evaluations of the groups (majority and minority) after 320 trials as a function of the type of the more frequent behavior (positive or negative). Error bars represent the standard error of the mean.

Figure 2

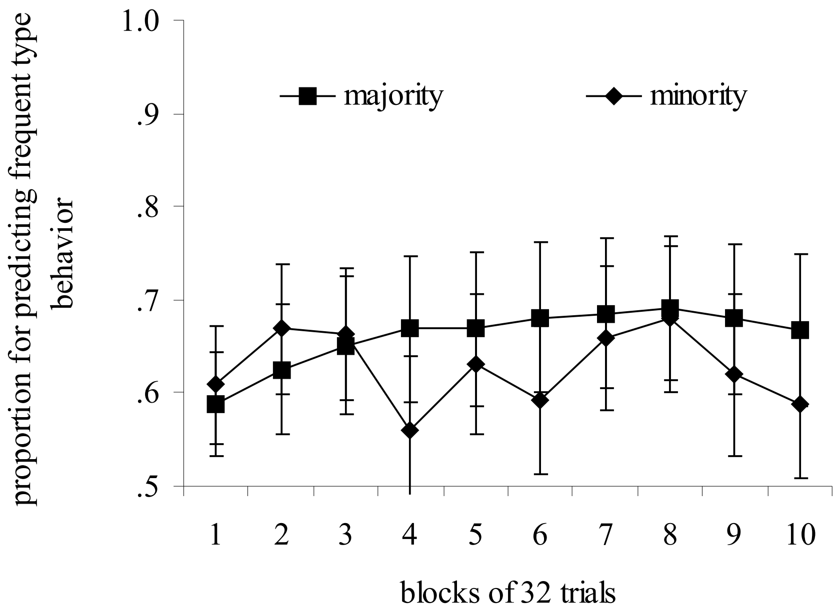

Figure 2. Mean conditional choice proportions for the more frequently presented behavior as a function of the type of group (majority or minority) on the trial. Error bars represent the standard error of the mean.

Figure 3

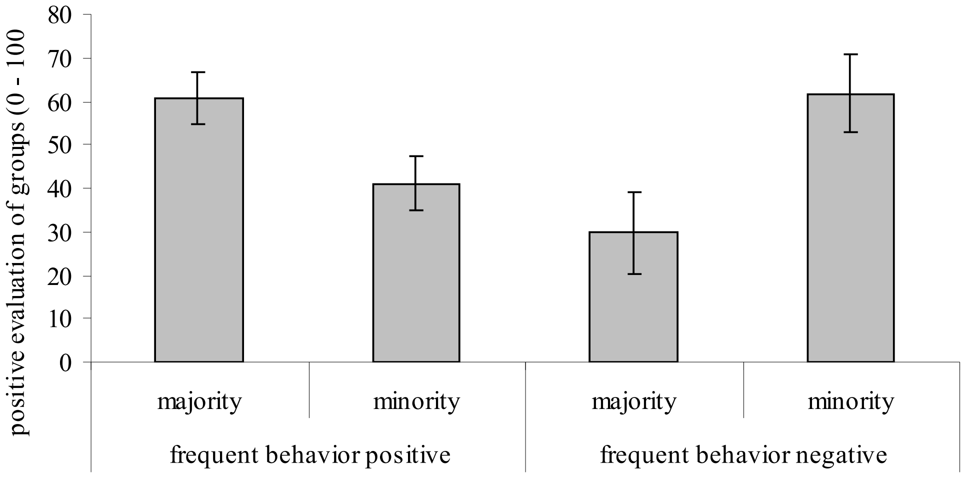

Figure 3. Mean evaluations of the groups (majority and minority) after 320 of not monetarily reinforced predictions as a function of the type of the more frequent behavior (positive or negative). Error bars represent the standard error of the mean.

Figure 4

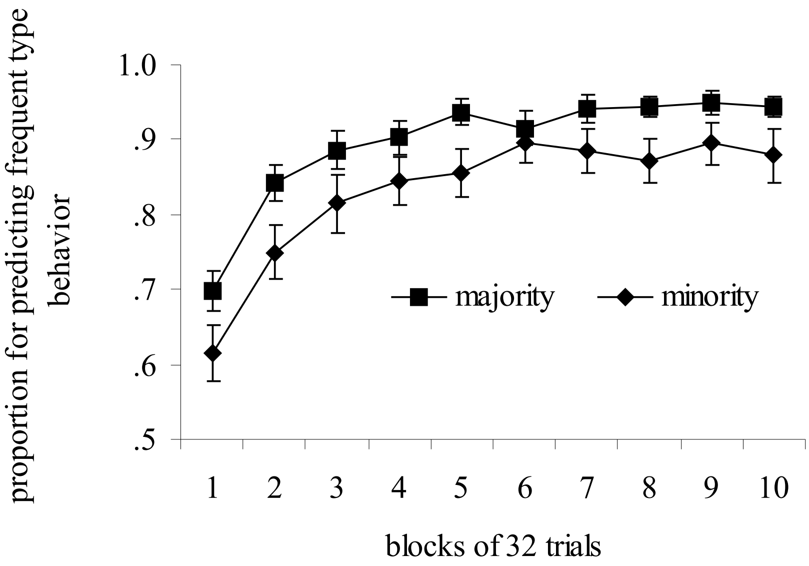

Figure 4. Mean conditional choice proportions for the more frequently presented behavior as a function of the type of group (majority or minority) on the trial. Error bars represent the standard error of the mean.

Figure 5

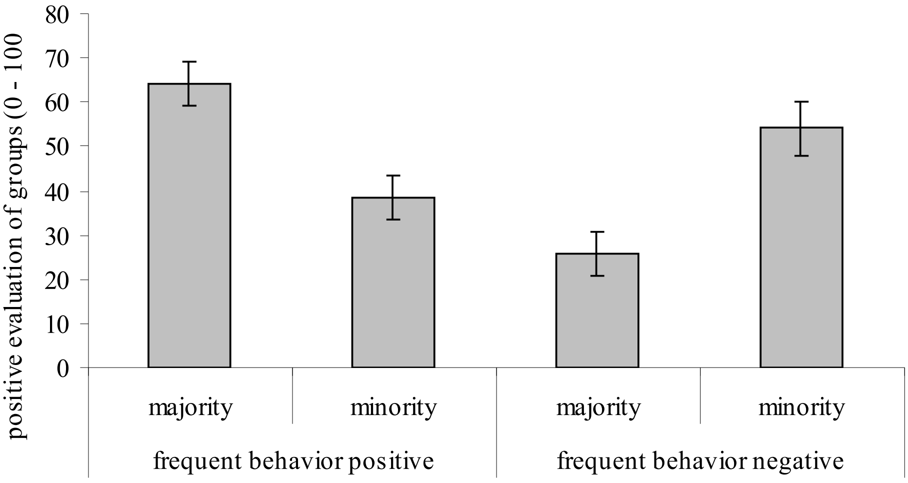

Figure 5. Mean evaluations of the groups (majority and minority) after 320 predictions as a function of the type of the more frequent behavior (positive or negative). Error bars represent the standard error of the mean.

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