

How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to Stereotype-Associated Dimensions

Deutsch, Roland; Fazio, Russell H.

Postprint / Postprint

Zeitschriftenartikel / journal article

Zur Verfügung gestellt in Kooperation mit / provided in cooperation with:

www.peerproject.eu

Empfohlene Zitierung / Suggested Citation:

Deutsch, R., & Fazio, R. H. (2008). How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to Stereotype-Associated Dimensions. *Journal of Experimental Social Psychology*, 44(4), 1020-1034. <https://doi.org/10.1016/j.jesp.2008.03.001>

Nutzungsbedingungen:

Dieser Text wird unter dem "PEER Licence Agreement zur Verfügung" gestellt. Nähere Auskünfte zum PEER-Projekt finden Sie hier: <http://www.peerproject.eu> Gewährt wird ein nicht exklusives, nicht übertragbares, persönliches und beschränktes Recht auf Nutzung dieses Dokuments. Dieses Dokument ist ausschließlich für den persönlichen, nicht-kommerziellen Gebrauch bestimmt. Auf sämtlichen Kopien dieses Dokuments müssen alle Urheberrechtshinweise und sonstigen Hinweise auf gesetzlichen Schutz beibehalten werden. Sie dürfen dieses Dokument nicht in irgendeiner Weise abändern, noch dürfen Sie dieses Dokument für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen.

Mit der Verwendung dieses Dokuments erkennen Sie die Nutzungsbedingungen an.

gesis
Leibniz-Institut
für Sozialwissenschaften

Terms of use:

This document is made available under the "PEER Licence Agreement". For more Information regarding the PEER-project see: <http://www.peerproject.eu> This document is solely intended for your personal, non-commercial use. All of the copies of this documents must retain all copyright information and other information regarding legal protection. You are not allowed to alter this document in any way, to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public.

By using this particular document, you accept the above-stated conditions of use.

Mitglied der

Leibniz-Gemeinschaft

Accepted Manuscript

How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to Stereotype-Associated Dimensions

Roland Deutsch, Russell H. Fazio

PII: S0022-1031(08)00034-6

DOI: [10.1016/j.jesp.2008.03.001](https://doi.org/10.1016/j.jesp.2008.03.001)

Reference: YJESP 2064

To appear in: *Journal of Experimental Social Psychology*

Received Date: 10 August 2007

Revised Date: 28 February 2008

Accepted Date: 7 March 2008



Please cite this article as: Deutsch, R., Fazio, R.H., How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to Stereotype-Associated Dimensions, *Journal of Experimental Social Psychology* (2008), doi: [10.1016/j.jesp.2008.03.001](https://doi.org/10.1016/j.jesp.2008.03.001)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Running head: PERCEPTION AND SUBTYPING

How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to
Stereotype-Associated Dimensions

Roland Deutsch
Ohio State University
and
Russell H. Fazio
Ohio State University

Please address correspondence concerning this article to:

Dr. Roland Deutsch
Lehrstuhl fuer Psychologie 2
Universitaet Wuerzburg
Roentgenring 10
97070 Wuerzburg, Germany

Email: deutsch@psychologie.uni-wuerzburg.de

Abstract

Two experiments examined the relation between stereotype disconfirmation and attentional processes. Using an instrumental learning-paradigm, we successfully simulated stereotype acquisition and the subsequent subtyping of disconfirming exemplars. While replicating established markers of subtyping, the present research demonstrates a hitherto neglected cognitive consequence of subtyping: Predictable stereotype disconfirmation increased attention to features that facilitated discriminating between confirming and disconfirming exemplars, and reduced attention to features associated with the original stereotype. These effects were not observed when stereotype disconfirmation was not easily predictable and, hence, subtyping proved difficult. The discussion focuses on implications for research on subtyping and stereotype change.

How Subtyping Shapes Perception: Predictable Exceptions to the Rule Reduce Attention to Stereotype-Associated Dimensions

As other concepts in memory, social stereotypes play an important role in helping humans to navigate through their social environment. To fulfill this function, they must be accurate to some degree (Judd & Park, 1993). Yet, negative stereotypes about social groups tend to persist even if they are fairly inaccurate. One important reason for the perseverance of inaccurate negative stereotypes presumably is the limited contact between people who hold stereotypes and the members of stereotyped groups (Dovidio, Gaertner, & Kawakami, 2003). External factors, such as spatial or cultural segregation of groups, as well as internal factors such as negative expectancies and resulting avoidance behavior (Fazio, Eiser, & Shook, 2004; Towles-Schwen & Fazio, 2003) contribute to limited contact, and prevent the experience of stereotype disconfirmation.

But even with sufficient contact, a correction of the stereotype may not ensue. One reason why stereotypes persist in the face of extended inter-group contact is the process of subtyping (Brewer, Dull, & Lui, 1981; Taylor, 1981). In its course, individuals who disconfirm the stereotype are grouped into a new subcategory that is mentally segregated from the rest of the group, thereby leaving the stereotype intact. For example, encountering a lawyer who is very introverted may lead people to conclude that she is a very atypical lawyer, and therefore not representative for the group as a whole. The introvert lawyer may then be put in a new subordinate category (Kunda & Oleson, 1995). Because disconfirming exemplars are excluded from the superordinate category, subtyping maintains or even reduces the perceived variability of the stereotyped group (e.g., Maurer, Park, & Rothbart, 1995). Moreover, it maintains or makes

the average of the stereotype more extreme (e.g., Hewstone, Macrae, Griffiths, Milne, & Brown, 1994).

Stereotype-disconfirmation is more likely to increase the perceived variability of the category and change its central tendency if there is little opportunity for subtyping. Research has established a number of preconditions of subtyping (Richards & Hewstone, 2001). For instance, subtyping is more likely to occur when the disconfirming exemplars deviate from the stereotype in an extreme (e.g., Kunda & Oleson, 1997) or atypical (e.g., Weber & Crocker, 1983) manner. For example, subtyping of a lawyer would be more likely if he was extremely vs. moderately introverted. Moreover, subtyping is more likely to occur if the disconfirming exemplars have some salient discriminative attributes (e.g., Kunda & Oleson, 1995). For example, an introverted lawyer would be more likely to be subtyped if, besides being introverted, he would have a particular style of clothing that distinguishes him from typical lawyers. Finally, research indicates that subtyping is more likely to occur if the social perceiver has relatively high cognitive capacity (e.g., Yzerbyt, Coull, & Rocher, 1999).¹

Why do people engage in subtyping? As Richards and Hewstone (2001) argue, “much of the literature on subtyping either explicitly or implicitly suggests that perceivers are motivated to maintain rather than to change their stereotypes in the face of disconfirming information” (p. 56). The specific motives for this conservatism, however, may be quite diverse. For example, members of advantaged groups may be motivated to maintain negative stereotypes about stigmatized minorities “...because they use the stereotypes to justify their social order, their sense of superiority to others, or their own behavior” (Kunda & Oleson, 1995, p. 566). In addition to social motives of this kind, other motives may also be regarded as causes of fencing off stereotype-disconfirming observations. Revising one’s well established and simple

stereotypes may collide with a need for simplicity and cognitive closure (Kruglanski & Freund, 1983), the need to think and behave consistently (Festinger, 1957), or the need to maintain positive self-views, such as being an unbiased person (Olson & Fazio, 2004).

But even without a direct and specific motivation to maintain the stereotype, the operation of more general mechanisms may prevent change. For example, being confronted with a person who disconfirms stereotypes may evoke surprise, which then triggers a search for specific reasons why this particular person deviates from the stereotype (Kunda & Oleson, 1995). Finding such person-specific reasons, however, then may protect the general stereotype from change. For example, finding that the introverted lawyer you just met is working full time for a human rights group may serve as a sufficient explanation for why he is different. At the same time, human rights activists may be perceived so different from typical lawyers that the particular lawyer is no longer seen as representative of the group of lawyers as a whole. Hence, in this case, the desire to explain unexpected findings triggers processes that may result in stereotype preservation. Also, subtyping may be an indirect consequence of action-control in general (e.g., Hommel, Müsseler, Aschersleben, & Prinz, 2001). Social interactions are more or less rewarding not only depending on the interaction partners per se, but also depending on the specific behaviors performed during the interaction with those partners. Stereotypes may provide their holders with scripts, specifying how to interact with members of stereotyped groups. If single members, however, deviate predictably from the stereotype in an action relevant manner, cognitively separating them from the rest is a functional means for optimizing behavior and the ensuing outcomes. In line with this notion, Shepard (1987) proposed that objects with similar consequences are grouped together in memory. Consequently, disconfirming exemplars may be subtyped, because they yield different outcomes than the average category member.

While previous research has primarily studied the consequences of subtyping for the structure and stability of an existing stereotype, the present research investigates how subtyping may change people's attention to and usage of perceptual dimensions related to the stereotype and subtype. Typically, social stereotypes are based on perceivable features of individuals, which signify a group membership. For example, basic social categories such as gender, ethnic origin, and age are relatively clearly discriminable based on perceptual dimensions such as skin tone, skin structure or body shape. The same is true for other social categories, albeit the relation between perceptual features and group membership may be fuzzier. Examples are dress codes associated with certain professions, or verbal accents that may correlate with socioeconomic status. Similarly, specific perceivable features typically are related to subtypes. For example, Black businessmen sometimes have been considered as representing a subtype of the group of Blacks in general (e.g., Kunda & Oleson, 1995). In this case, subtype membership can be inferred from features relating to the ethnic descent, and features relating to the profession, such as wearing a business suit. It is the latter features that particularly help distinguishing subtypes from regular types.

As diverse as the mechanisms underlying subtyping may be, they imply similar changes in attention. According to the mechanism based on action-control, actors aim to discriminate between subtypes and regular types because they require different behaviors if one is to achieve favorable (or not unfavorable) outcomes. People should actively search for (and choose) the appropriate behaviors for a rewarding social interaction with the different targets. Consequently, attention towards those features that help distinguishing subtypes from regular types should be generally increased. Also, as a consequence of this need to match behavior to the target, attention to the major stereotype-related dimension should be reduced because it alone is not sufficient for

action control. According to the mechanism based on expectancy-violation induced surprise (Kunda & Oleson, 1995), people are motivated to explain the violation, thereby potentially avoiding future surprises. To achieve this goal, focusing on the specific features that discriminate subtypes from regular types is a functional strategy. Similar to these non-directional motives, motives directed at protecting the stereotype from change may increase people's attention to and usage of features that help identifying subtypes. If motives such as a need to justify social hierarchies, a need to view oneself as being consistent or acting according to norms of fairness fuel subtyping, it is essential for stereotype holders to recognize exemplars belonging to the subtyped category in an efficient way. Also, atypical features of the subtyped exemplars should be processed with priority, because they provide the justification for dismissing these exemplars as evidence against the validity of the stereotype.

In essence, although subtyping can be driven by different motives and cognitive mechanisms, it may uniformly shape attention to and usage of perceptual dimensions related to the stereotype and subtype. In other words, subtyping produces a change in the mental representation of the original category. In particular, we hypothesize that subtyping increases attention to certain perceptual dimensions, those that facilitate discriminating between the subtype and the superordinate category, such as dress in the case of the Black businessman subtype. Likewise, we hypothesize that subtyping decreases attention to perceptual dimensions that were originally thought to be predictive of the superordinate category, such as skin tone in the case of racial stereotypes. To our knowledge such attentional changes as a consequence of subtyping have not been demonstrated. Research on subtyping has emphasized that the content and valence of the original stereotype remain intact, despite the experience of disconfirming instances. The present research does not question the validity or significance of these effects.

However, the novel possibility we wish to consider is that subtyping does produce changes in the attention given to the dimensions of relevance. In this way, subtyping may lead to changes in the mental representation defining the original category.

The Present Research

In the present research, we designed two experiments to test the hypothesized effects of subtyping on perception. The nature of our hypotheses implied two deviations from traditional research paradigms, which we will explain in the following paragraphs: (a) The use of artificial stimuli and stereotypes, and (b) the use of positive stereotypes.

Traditionally, researchers have relied on pre-existing stereotypes and presented research participants with disconfirming exemplars. Different from this general approach, we studied stereotypes about artificial instead of naturalistic stimuli, as has been successfully done in a few previous studies (e.g., Rothbart & Lewis, 1988). We did so in order to maximize experimental control over the perceptual features underlying the stereotype and the subtype. Maximal control over these features is desirable because our hypotheses address the attention that these features receive as a function of subtyping. To achieve this, we simulated the acquisition of stereotypes and subtypes using the paradigm BeanFest, which was recently developed to study attitude-learning as a function of exploratory behavior (Fazio et al., 2004).

BeanFest is a computer game simulating a world that consists of beans of different types. The participant's goal is to gain points by repeatedly deciding which specific beans to approach and which beans to avoid. Each bean has a positive or negative value, and participants lose points when they approach a negative bean, whereas they win points when they approach a positive bean. Participants also can reject beans, in which case the participant's point value is unaffected. However, in such cases, the value of the bean (i.e., the points that would have been

won or lost if the bean had been approached) is not learned. At any given time, the participant's cumulative point value ranges from 0 to 100. The beans themselves differ by shape and number of speckles. They can be viewed as forming a 10 by 10 matrix in which one dimension represents the shape of the bean, ranging from circular to oval to oblong, and the other dimension represents the number of speckles, ranging from one to ten (see *Figure 1*).

Within the matrix, 24 beans were selected to represent a very simple reward scheme. For example, in *Figure 2 A*, the 12 beans with very few speckles were bad, i.e., they yielded a loss when approached, whereas the 12 beans with very many speckles were good, i.e., they yielded a gain when approached. In both experiments, we first induced a stereotype about the positive area by handing out written information, purportedly stemming from other participants who had played the game earlier. This method had been successfully used to simulate socially transmitted stereotypes or prejudice in previous research (Fazio et al., 2004). We decided to use only positive stereotypes in the present research, mainly because previous research indicates that participants vigorously avoid beans presumed to be negative (Fazio et al., 2004). Therefore, in the standard set-up, negative stereotypes would make stereotype disconfirmation very unlikely, simply because participants can fully avoid the purportedly negative beans. With a positive stereotype, however, disconfirmation will occur frequently, because the positive expectation will promote approach behavior (Fazio et al., 2004). Thus, the present operationalization focuses on positive stereotypes in order to increase the likelihood that participants will experience exceptions. Nevertheless, we are not aware of any a priori reasons why the effects of disconfirmation, once experienced, should differ for positive and negative stereotypes in any qualitative manner.

Hence, with respect to *Figure 2*, the stereotype-inducing reports presented to participants suggested that the beans with many speckles were good. We manipulated the way in which this

stereotype was disconfirmed later during the game. As a control group, we had one condition without any disconfirmation at all (*Figure 2 B*). Those in another condition worked with a matrix in which atypical, speckled beans were of negative valence, whereas the more typical speckled bean was positive. In this case, the disconfirmation was clustered around very oblong beans with many speckles (*Figure 2 C*). Clustered disconfirmation promotes subtyping, because it associates a salient perceptual feature with stereotype disconfirmation (Queller & Smith, 2002), making it easier to abstract disconfirming exemplars into a subtype (Richards & Hewstone, 2001). A real-world example of clustered disconfirmation would be meeting a number of introverted lawyers, who are also casually dressed and are working for human rights organizations. Finally, both experiments included a group with a matrix that included the same degree of disconfirmation as the clustered disconfirmation matrix, but this time distributed widely across the category of beans with many speckles (*Figure 2 D*). That is, speckled beans were sometimes negative, irrespective of their overall shape. In this way, we could establish stereotype-disconfirmation without an opportunity to identify an atypical subtype -- a condition well known to lead to stereotype change (Richards & Hewstone, 2001). In terms of our running example regarding lawyers, distributed disconfirmation would involve meeting a number of introverted lawyers who were dressed diversely and who worked for a variety of organizations. Note that the average degree of disconfirmation was identical in the clustered and distributed group.

In both experiments, BeanFest was used as a learning paradigm, involving reward contingencies similar to those described above, and attention towards stereotype-relevant and stereotype-irrelevant perceptual dimensions was assessed using multidimensional scaling (*MDS*; see Nosofsky, 1992). *MDS* is a multivariate technique, which is frequently used by cognitive scientists to study attention and categorization.² An *MDS* analysis uses similarity judgments

between objects as input, and returns a number of dimensions, which constitute a multidimensional space in which the distance between any two stimuli reflects the extent to which they have been rated as dissimilar. In the present study, the MDS is conducted to analyze how many perceptual dimensions people use to judge the similarity of beans. The most obvious expectation is that people base their judgments on two dimensions, namely shape and speckles. Besides a number of dimensions, the MDS returns stimulus weights, which indicate the location of objects in the multidimensional space. Objects are located in such a way that the distances between the objects reflect their perceptual similarity. In the present experiment, this would mean that each bean has individual dimension weights (or coordinates) that determine its place within the hypothesized dimensions of shape and speckles. Beans that are close together in this space were, on average, judged to be more similar than beans that are far apart.

In addition, and most relevant for our purposes, some variants of the MDS also return dimension weights for each participant. These dimension weights give an estimate of the extent to which the individual weighted each dimension in making his or her similarity ratings. Applied to the present experiments, this means that some participants may rely more on shape when judging the similarity of beans, whereas others would more rely on the number of speckles, whereas other participants might give equal weight to the two dimensions. In other words, people may attend more or less to a given dimension. In the MDS, such differences would be mirrored in individual differences in the weights associated with the extracted dimensions that correspond to shape and speckles. The MDS returns a number of dimensions, but their psychological meaning (e.g., whether they represent shape or speckles), however, must be inferred from the multidimensional space. To this end, one can correlate stimulus features (e.g., the number of speckles) with stimulus weights of a given dimension. If the correlation is sufficiently high, one

can infer that a given dimension represents the given stimulus feature. This is comparable to a factor analysis, where the psychological meaning of single factors must be inferred from the factor loadings of single items.

Experiment 1 was geared towards a first test of the major hypotheses regarding the differential impact of subtyping on attention-allocation and behavioral performance. In Experiment 2, we sought not only to replicate the effects observed in Experiment 1 with a different stereotype, but also to establish that the present paradigm replicates known consequences of subtyping, despite the paradigm's reliance on artificial stimuli.

Experiment 1

In Experiment 1, we aimed to test the following hypotheses derived from our theoretical framework. First, because the subtype in the clustered disconfirmation condition was based on a combination of the stereotypic dimension (number of speckles) and the other dimension (shape), we expected subtyping to increase attention towards the non-stereotypic dimension (shape), whereas no such effect should occur with distributed disconfirmation. Second, we expected a reduced attention towards the stereotype-related dimension (speckles) with subtyping, but not with distributed disconfirmation. With clustered disconfirmation, people learn that they cannot predict the effect of a bean from speckles alone, and therefore are expected to pay relatively less attention to it. It is important to note that pre-tests had revealed that shape is perceptually dominant over speckles. More specifically, MDS analyses indicated that about 60% more weight was given to the shape of the beans than to their number of speckles when similarity judgments were sampled without having played the BeanFest game. Hence, any effect of experiencing stereotype disconfirmation must be judged against this baseline. Consequently, we expected the dominance of shape over speckles to be strongly reduced with no disconfirmation or with

distributed disconfirmation, but less so with clustered disconfirmation. In this latter case, participants' experiences will give them more reason to attend to shape. Finally, we anticipated that with clustered disconfirmation, people would behave more successfully than with distributed disconfirmation. Mentally separating the disconfirming exemplars from the rest, as well as the enhanced predictability of disconfirmation makes solving the task much easier with clustered than with distributed disconfirmation.

Method

Participants and Design

Seventy-one Ohio State University students enrolled in introductory psychology courses (43 females and 28 males) participated in this experiment for research credit. At most, four participants were present for each session. Data from one participant were excluded from the analyses because she was not a native speaker and indicated that she did not understand the instructions correctly. Data from another participant were excluded because he already had participated in a very similar experiment using the same paradigm. Hence, the following analyses are based on $N = 69$. The design involved three experimental groups, representing different levels of stereotype-disconfirmation: no disconfirmation, clustered disconfirmation, and distributed disconfirmation of the stereotype (see *Figure 2 B, C, and D*, respectively).

Materials

All stimuli were derived from the 10 by 10 bean matrix depicted in Figures 1 and 2 A. For the learning phase, 12 beans from the first and second columns (i.e., beans with few speckles) and 12 beans from the 9th and 10th columns (i.e., beans with many speckles) were selected. For the similarity ratings, 12 beans were selected in a way as to give a representative sample of the bean matrix. Referring to the 10 by 10 coordinate system, the following beans

were chosen: 1 by 1, 1 by 10, 2 by 2, 2 by 5, 2 by 9, 5 by 1, 5 by 10, 9 by 2, 9 by 6, 9 by 9, 10 by 1, and 10 by 10. These beans were used to generate 66 pairs of beans, representing each possible combination of non-identical beans.

Procedure

When participants arrived at the lab, they were seated in individual cubicles and provided extended written instructions for BeanFest. The experimenter read the instructions aloud, while the participants read along.

General instructions. Participants were informed that they would play a game that involves beans, which they could choose to approach or to avoid. Every bean would have a positive or negative point value throughout the game. They could learn about the value of a given bean by approaching it, but not if they would avoid it. Thus, in order to learn which beans to choose and which beans to avoid, they would first have to sample a few. They were told that their task was to gain as many points as possible, and that reaching 100 points represented winning the game, whereas reaching 0 points represented losing the game. After having won or lost, the game would be re-started. Because pilot-testing had revealed that some people had difficulty discerning the critical dimensions of the stimuli in the number of trials that were to be presented, these instructions explicitly informed the participants that the beans varied in shape (“from circular to oval to oblong”) and in the extent to which they are speckled (“marked by anywhere from very few to some to many speckles”). In addition, the information that was to be displayed on the computer screen was illustrated and described. Finally, the instructions indicated that the experiment would begin with a practice block of 6 trials. Participants were forewarned that the 6 beans to be presented would be just a few of the ones that they would see during the actual experiment, but that the beans would have the same value as they would during

the experiment. These trials were described as their first opportunity to begin to learn about some of the beans and to familiarize themselves with the feedback displays. They were explicitly instructed that, given these goals, they should respond YES (i.e., *approach*) on each of the 6 practice trials.

Stereotype induction. After the general instructions, but prior to the beginning of play, stereotypic knowledge about some of the beans was conveyed. They were told that the focus of the BeanFest experiment was not on individual learning, but on how people learn across generations (see Fazio et al, 2004; Experiment 5 for a more detailed description of the method). They were presumably later generation participants. Earlier generations of participants in the game had provided written observations and suggestions, and they would have access to two such generation reports. Under this guise, participants received suggestions asserting that beans with many speckles are good, but that nothing clear could be said about the valence of the other beans. Hence, the information conveyed positive expectations about the rightmost area of the bean-matrix, which were then confirmed or disconfirmed to varying degrees, depending on experimental conditions.

Game phase. At the beginning of BeanFest, participants underwent a practice block of six trials. Three beans from each of the two regions the matrix was presented, all of which confirmed the conveyed expectations. Participants were asked to accept each practice bean, in order to familiarize themselves with the feedback and point displays and begin to associate a few specific beans with their point values. When finished with the practice phase, participants started the actual game, which consisted of four blocks of 24 trials. The 24 trials involved the beans within the selected regions of the matrix (see *Figure 2 A-D*). Each bean was presented once in each block in a random order; thus, all 24 beans were seen four times. During a trial, participants were

presented with a bean in the upper portion of the monitor. They had to indicate whether they wanted to accept or reject the bean. Participants responded by pressing either the “yes” or the “no” button on the keyboard. After responding to each bean, the lower portion of the monitor was adjusted according to the participant’s decision. All of the information about the participant’s point value was located in the lower right corner of the monitor. The point value was represented both numerically and graphically as a bar ranging from 0 to 100. These fluctuated in response to the participant’s decision to accept a bean as a function of the bean’s value. In the lower left corner of the monitor, participants were presented with information about their response and the bean’s value. The participant’s response appeared as either *yes* or *no*. The bean’s value appeared below the response, but only if the participant chose to approach the bean. Participants started the game with 50 points, with reaching 100 representing winning the game, and reaching 0 representing losing the game. If participants won or lost, the game restarted. Participants would restart at 50 points. The game restarted as many times as the participants won or lost. With any restarted games, the beans retained their original values. Thus, participants did not have to relearn the beans if they played multiple games.

Similarity judgments. When all participants were finished playing BeanFest, they received instructions on the computer screen for the similarity judgments, which served as input for the MDS-analyses. All 66 pairs were presented in a sequence that was randomly determined for each participant. For each pair, participants were instructed to judge their perceptual similarity on 9-point scales labeled from *not at all similar* to *very similar*. They were instructed to give quick responses, but were not paced. Before the first trial, participants saw an overview of the bean-matrix similar to Figure 1 for 15 seconds, to give them a frame of reference for judging perceptual similarity. At the beginning of each trial, a pair of beans was presented on the

screen, with one bean on the left and one on the right. After the judgment for a bean was collected, the next trial started.

Results

The results will be organized along two questions. First, we will examine how learning about clustered and distributed disconfirmation influenced attention towards the two perceptual dimensions of the bean matrix. Second, we will examine how our experimental manipulation affected participants' performance during the game.

Attentional Effects

We predicted that the type of stereotype disconfirmation would affect which features of the exemplars receive more or less attention. To test this prediction, we conducted MDS analyses using the perceptual similarity ratings as input. We used the ALSCAL module that is part of SPSS 14.0 for the following analyses. We performed an INDSCAL analysis, which has the advantage of returning not only stimulus weights for a given similarity space, but also individual weights, which provide an estimate of how much a person is making use of the given dimension. We then analyzed the individual dimension weights as a function of experimental conditions. This way, we inferred how much attention was given to the dimensions as a function of clustered vs. distributed disconfirmation.

Because the bean-matrix involves two dimensions (shape and speckles), it is plausible to expect that the MDS would return the best fit for a two-dimensional solution. In this (and the following) Experiment, this was actually not the case. Instead, a three-dimensional solution provided a better, and overall acceptable fit, $Stress = .14$, $RSQ = .82$. We suspected that the third dimension might be related to the novelty of a given bean, because some of the beans that were used for the similarity ratings had not been presented during the game. To interpret the meaning

of a given MDS dimension, we predicted the stimulus weight of each bean on this dimension (i.e. a bean's location on this dimension) from each bean's *number of speckles*, *shape*, and *novelty*. Particularly, these three features were entered into a regression analysis as simultaneous predictors. Speckles and shape were coded ranging from 1 to 10 according to their location in the 10 X 10 bean matrix. Novelty was coded as a dichotomy, reflecting whether the bean had or had not been presented during the game. This analysis revealed that shape was related to Dimension 1, speckles were related to Dimension 2, and novelty was related to Dimension 3 (see Table 1). Because the novelty dimension is not relevant for the theoretical predictions, further analyses will be focused on shape and speckles. Also, novelty had an incremental *RSQ* of only .081, indicating that this dimension was relatively unimportant for participant's perceptual space.

How did the presence and distribution of disconfirming events affect people's perceptions of the beans? To address this question, we examined each participant's squared dimension weights for the dimensions identified as shape and speckles. Note that whereas shape and speckles are objective stimulus-features, the individual dimension weights serve as estimates of the extent to which any given participant relied on shape versus speckles when judging similarity. In Figure 3, the mean squared dimension weights are plotted as a function of type of disconfirmation. Participants in the condition with no disconfirmation and with distributed disconfirmation gave relatively equal weight to the stereotypic dimension (speckles) and the non-stereotypic dimension (shape).⁴ With clustered disconfirmation, however, participants gave considerably more weight to the non-stereotypic dimension compared to the stereotypic dimension (see Figure 3). This interpretation is supported by a 2 (dimension) X 3 (disconfirmation) ANOVA on the squared dimension weights. This analysis revealed a main effect of dimension, indicating that generally more weight was given to shape, $F(1, 66) = 9.45, p$

$= .002$, $\eta^2 = .13$. Most importantly, however, this main effect was qualified by an interaction of dimension and disconfirmation, $F(2, 66) = 4.66$, $p = .013$, $\eta^2 = .12$. Simple contrasts revealed that different weight was given to shape and speckles with clustered disconfirmation ($p < .001$), but not with distributed ($p = .988$) and no disconfirmation ($p = .196$). Furthermore, participants in the clustered disconfirmation condition weighted shape more heavily than did those in the distributed condition ($p = .009$) and the no disconfirmation condition ($p = .087$). Likewise, participants in the clustered disconfirmation condition weighted speckles less heavily than did those in the distributed condition ($p = .008$) and the no disconfirmation condition ($p = .057$). Thus, without disconfirmation and with distributed disconfirmation, shape and speckles were weighted equally, whereas with clustered disconfirmation, shape received significantly more weight, and speckles received significantly less weight in judgments of perceptual similarity.⁵

Learning

To validate our presumption regarding the ease of learning cases of clustered versus distributed disconfirmation, we examined participants' game performance. We first recoded participants' responses as to whether they were correct (approaching positive beans, avoiding negative beans) or incorrect (approaching negative beans, avoiding positive beans). We then submitted the overall percentage of correct responses to a 4 (Block) X 3 (Disconfirmation) ANOVA, with the first factor varying within, the last factor varying between subjects (see Figure 4). Because percentages violate homogeneity assumptions of the ANOVA, we performed an arcsine transformation (Kirk, 1968) on the raw percentages before submitting them to the ANOVA. For the sake of readability, we report the raw percentages in Figure 4. The analysis revealed that participants' performance significantly increased with the number of blocks, $F(3, 198) = 23.39$, $p < .001$, $\eta^2 = .26$. Participants improved from 81% correct responses in the first

block to 92% correct responses in the fourth block. Most important to the present hypotheses, however, people were differentially successful depending on the type of disconfirmation, $F(2, 66) = 54.83, p < .001, \eta^2 = .62$.⁶

Contrast analyses revealed that participants performed best without disconfirmation ($M = 96\%$, $SEM = 1.53$), second best with clustered disconfirmation ($M = 88\%$, $SEM = 1.53$), and worst with distributed disconfirmation ($M = 79\%$, $SEM = 1.44$) with distributed disconfirmation. This rank order was apparent from block one to block four (all $p < .05$). In essence, as anticipated, participants were more successful with clustered than with distributed disconfirmation.

Discussion

The results of Experiment 1 proved supportive of our hypotheses. To study disconfirmation of a pre-existing stereotype, we induced expectations within the artificial world of BeanFest. Drawing upon a technique previously used by Fazio et al. (2004), we provided participants with reports purportedly stemming from previous generations of participants, indicating that beans with many speckles were positive. This induction was successful, as participants performance during the game was very good from block 1 on, indicating that they followed the by-and-large correct rule conveyed in the reports. However, two of the three groups experienced occasional disconfirmation of their positive expectations during the game.

We predicted that the opportunity to mentally cluster exceptions to the rule would have consequences for individuals' mental representations of the original category. More specifically, we predicted that subtyping would affect the amount of attention that individuals would subsequently direct towards those features relevant to the stereotype and those features relevant to the exceptions. To test this prediction, we examined participants' judgments of the perceptual

similarity of pairs of beans following the BeanFest game. In particular, we studied the weight that participants subsequently gave to the stereotype relevant (speckles) and irrelevant (shape) features of beans, by using the similarity judgments as input to an individual difference MDS. In accord with the hypothesis, we observed that shape and speckles received equal attention with no disconfirmation or distributed disconfirmation. With clustered disconfirmation, however, attention towards the stereotype-relevant feature was reduced, whereas attention to the feature distinguishing exceptions from the rule was increased. These results suggest that while subtyping may have the consequence of stabilizing the original category, it also may make it less likely that people attend exclusively to the dimensions originally associated with the stereotype. For example, when learning about the subtype of Black businessmen, attention to features associated with ethnic descent (e.g., skin tone) may be reduced, whereas attention to features that help recognizing the subtype (e.g., signs of a particular profession, such as clothing) may be increased.

We also observed, as expected, that when the exceptions to the rule were clustered, i.e., when they were perceptually atypical exemplars of the class of speckled beans, participants better learned to avoid them as the game proceeded. This suggests that the opportunity to mentally group disconfirming exemplars increases one's ability to respond favorably to atypical exemplars, whereas omitting such an opportunity makes it harder to treat the exceptions to the rule in an appropriate way. In a real world example, this would suggest that forming a subtype of introverted lawyers would help people to engage in positive interactions with introverted vs. extraverted lawyers, because the subtype would be of help in discriminating between the two groups of lawyers.

Experiment 2

In Experiment 2, we wanted to replicate and extend the results obtained in Experiment 1. The overall operationalizations and hypotheses of Experiment 2 were similar to Experiment 1. However, to make sure that the judgmental weighting of dimensions observed in Experiment 1 was not driven by the particular perceptual features linked to the stereotype (i.e., speckles) and the subtype (i.e. oblong shape and speckles), Experiment 2 modified the stereotype and subtype that was operationalized within the BeanFest paradigm. Specifically, in Experiment 2 the positive stereotype was conveyed regarding round beans instead of speckled beans. Hence, in Experiment 2, the stereotype was related to the naturally dominant dimension of shape (see footnote 1). Consequently, we expected that the weight given to shape should be reduced with clustered compared to distributed and no disconfirmation. Conversely, we expected that the weight given to speckles should be enhanced with clustered disconfirmation compared to distributed and no disconfirmation. Furthermore, a different sampling of bean pairs was employed to obtain the similarity judgments than in Experiment 1. To get a more representative sample of the bean matrix, more beans from the center of the bean-matrix were sampled, while reducing the total number of beans by two in order to relieve the burden on the participants.

More importantly, Experiment 2 also aimed to provide direct evidence that our specific paradigm (i.e. clustered disconfirmation in BeanFest) indeed induces subtyping processes similar to those that have been observed with more naturalistic materials. Specifically, we aimed to test whether typical outcomes of subtyping (e.g., bolstering of the original stereotype) can also be observed in the present paradigm. To this end, additional dependent variables were introduced, particularly the belief in the validity of the stereotype and a measure of perceived category valence. If our paradigm maps processes similar to past subtyping research, clustered

disconfirmation should yield more positive estimates of category valence and a greater belief in the stereotype compared to distributed disconfirmation. Moreover, a de-categorization of disconfirming exemplars should occur with clustered but not with distributed disconfirmation.

To examine de-categorization, we introduced an explicit category label referring to the stereotypic beans. Specifically, for half of the participants, round beans were introduced as *Kambo* beans, and before they actually played the game, they practiced categorizing beans as *Kambo* vs. non-*Kambo*. After the game, we measured people's willingness to use the category *Kambo* in the same categorization task as before. Based on previous research on subtyping, we expected that clustered and, hence, predictable disconfirmation would result in de-categorizing the specific disconfirming exemplars, whereas such a specific effect should not occur for distributed and, hence, less predictable disconfirmation. Thus, subtyping on the basis of the distinct perceptual features that are associated with clustered disconfirmation was expected to decrease the likelihood that participants would continue to label the disconfirming exemplars as *Kambos*, even though such round beans were initially accepted as members of the *Kambo* category. Beyond that, we also wanted to explore how the presence vs. absence of a distinct category label might influence the attentional effects observed in Experiment 1. This is an interesting question because many stereotypes come with a ready to use explicit category label (such as *the elderly*, *Blacks* etc.), and one might argue that the presence of such strong labels may interfere with the attentional changes observed in Experiment 1.

Method

The methods of Experiment 1 and Experiment 2 were very similar. In the following, we describe only the unique aspects of Experiment 2 in detail.

Participants and Design

One-hundred twenty-seven Ohio State University students enrolled in introductory psychology (65 females and 62 males) participated in this experiment for research credit. The design was a mixed 2 (category learning present vs. absent) X 3 (disconfirming events: none vs. distributed vs. clustered) with both factors being manipulated between participants.

Materials

All stimuli were derived from the 10 by 10 bean matrix depicted in Figure 1. The stimuli selected for the learning phase, however, differed from Experiment 1. In particular, instead of using the two leftmost and two rightmost columns of the matrix, 12 stimuli from the upper two rows (i.e., round beans) and 12 stimuli from the lower two rows (i.e., oblong beans) were selected. Within the round beans, the three different reward patterns used in Experiment 1 were superimposed to the round beans, while the oblong beans were always negative. For the similarity ratings, 10 beans were selected in a way as to give a representative sample of the bean matrix, this time drawing more beans from the center area. Referring to the 10 by 10 coordinate system, the following beans were chosen: 1 by 1, 1 by 10, 2 by 5, 5 by 5, 5 by 8, 6 by 3, 6 by 6, 9 by 6, 10 by 1, and 10 by 10. These beans were used to generate 45 pairs of beans, representing each possible combination of non-identical beans.

Procedure

The procedure was similar to that of Experiment 1 with a few exceptions.

Category learning. Half of the participants began with a category-learning task. In this task, they were instructed that there was a specific category of beans, *Kambo* beans, about which they were going to learn in the future. To begin with, they were instructed that they should learn to discriminate Kambo beans from other beans. The only hint they got was that “Kambo beans

tend to be round”. Then they were successively presented with 20 round and 20 oblong beans, and instructed to indicate whether the given bean would be an exemplar of Kambo beans or not, using the appropriate key on the keyboard. Participants received error-feedback and showed nearly perfect performance after a few trials.

Stereotype induction. Stereotype induction was similar to Experiment 1, except that the verbally conveyed stereotype referred to the round beans as being good. In the category-learning group, the round beans were addressed as Kambo beans – the name the participants had just practiced.

Game phase. The game phase was identical to Experiment 1, except that the transposed reward scheme was used.

Similarity judgments. Collection of similarity judgments was identical to Experiment 1, except that only 10 beans were used, resulting in only 45 trials.

Re-classification: Those participants who had learned to classify round beans as Kambo beans again engaged in a classification task. They were presented with half of the bean matrix (i.e. 50 beans), sampled in a “checkerboard” manner, to reduce the working load of participants. Referring to the 10 by 10 coordinate system, beans were chosen according to the following schema: 1 by 2, 1 by 4, 1 by 6, 1 by 8, 1 by 10 for the first row; 2 by 1, 2 by 3, 2 by 5, 2 by 7, 2 by 9 for the second row, and so forth. All participants who worked on the re-classification task were presented with the same selection of stimuli. As a consequence of this selection, only two of the three disconfirming exemplars, and only four of the nine confirming exemplars were actually presented during the re-classification task. Participants were instructed to remember what they had learned earlier, at the beginning of the experiment, about how Kambo beans

looked. No reference was made to the game or the stereotype. Error feedback was omitted this time.

Final questions. Finally, we asked participants about their general impression of how good the round beans (or Kambo beans, in the category learning condition) had been. This question was used to study any changes in perceived category valence as a function of learning. Additionally, we asked participants about how valid they perceived the conveyed stereotype to be.

Results

Attentional Effects

As in Experiment 1, we tested our hypotheses regarding the attentional effects of stereotype disconfirmation by submitting the similarity judgments to an MDS, using an INDSCAL model. Mirroring the results observed in Experiment 1, the three-dimensional solution was superior to the two-dimensional solution with respect to *Stress* (0.15) and *RSQ* (0.82). To interpret the dimension, we performed three regression analyses, each using the features *number of speckles*, *shape*, and *novelty* as simultaneously entered predictors. The weights of dimension one, dimension two, and dimension three served as the dependent variable in the three analyses respectively. Results indicate that dimension one represents shape, dimension three represents speckles, and that dimension two represents novelty (see Table 2). As in Experiment 1, we refrained from analyzing the novelty dimension, as it is not relevant to the present hypotheses. The novelty dimension presumably emerged as a consequence of the sampling method used.

How did the presence and distribution of disconfirming events affect people's perceptions of the beans? In Figure 5, the mean squared dimension weights for the dimensions

indicative of shape and speckles are plotted as a function of type of disconfirmation. Participants in the condition without disconfirmation and with distributed disconfirmation gave more weight to the stereotypic dimension (shape) than participants who experienced clustered disconfirmation. Generally speaking, participants with clustered disconfirmation gave the most weight to speckles (the non-stereotypic dimension), participants without disconfirmation gave the least weight to speckles, and participants with distributed disconfirmation fell in-between the former two. This interpretation is supported by a 2 (dimension) X 3 (disconfirmation) X 2 (category label present or absent) ANOVA on the squared dimension weights. This analysis revealed a main effect of dimension, indicating that generally more weight was given to shape, $F(1, 121) = 236.7, p < .001, \eta^2 = .66$. This main effect was qualified by an interaction of dimension and disconfirmation, $F(2, 121) = 3.92, p = .022, \eta^2 = .06$. Contrast analyses revealed a significant effect of type of disconfirmation on use of the shape dimension ($p = .023$) and a marginally significant effect on attention to speckles ($p = .066$). Shape was weighted equally given no disconfirmation and distributed disconfirmation ($p = .98$), but less so with clustered disconfirmation ($p < .05$ for both comparisons). Speckles were weighted less without disconfirmation than with clustered disconfirmation ($p = .02$), but the weight of speckles in the distributed disconfirmation did not significantly differ from the other two conditions (both $p > .20$). No other main effect or interaction reached statistical significance (all $p > .10$), indicating that the presence vs. absence of a verbal category did not affect how participants weighed shapes and speckles.⁷

Categorization

Subtyping suggests that disconfirming exemplars are excluded from the original category. To test whether this process actually occurred in our paradigm, we compared the

probability of categorizing the two presented disconfirming exemplars vs. the four presented non-disconfirming exemplars as Kambo beans for those participants who experienced one of the two types of disconfirmation. An inspection of the means (see Figure 6) suggests that participants who experienced clustered disconfirmation were less willing to categorize disconfirming than confirming exemplars as Kambo beans. For those who experienced distributed disconfirmation, there was only a very small difference in the likelihood of categorizing confirming and disconfirming exemplars as Kambo beans. A 2 (disconfirmation) \times 2 (exemplar status) ANOVA with the first factor varying between, the second factor varying within participants supports this interpretation. As in Experiment 1, we performed an arcsine transformation (Kirk, 1968) on the raw probabilities before submitting them to the ANOVA. For the sake of readability, we report the raw probabilities in Figure 6. Overall, disconfirming exemplars were less likely to be categorized as Kambos than confirming exemplars, $F(1, 41) = 14.98, p < .001, \eta^2 = .27$. This main effect, however, was qualified by a two-way interaction of type of disconfirmation and exemplar status, $F(1, 41) = 4.43, p = .041, \eta^2 = .10$. Simple contrasts revealed that for disconfirming exemplars, categorization-likelihoods were lower with clustered than with distributed disconfirmation ($p = .038$), but not for confirming exemplars ($p = .780$). Moreover, with clustered disconfirmation, disconfirming exemplars were less likely categorized as Kambos than confirming exemplars ($p < .001$), whereas this difference was not significant with distributed disconfirmation ($p = .224$). The main effect of type of disconfirmation was marginally significant, $F(1, 41) = 3.53, p = .067$. In sum, this analysis suggests that participants with clustered disconfirmation, but not with distributed disconfirmation, selectively de-categorized disconfirming exemplars. What clearly had been a Kambo, during the early

category-learning portion of the experiment, was less likely to be considered a Kambo after subtyping of the beans that produced an unexpected negative outcome.

Valence of Round Beans

How did experiencing disconfirmation change participants' evaluation of the stereotype related set of stimuli as a whole? Participants who learned the category Kambo were asked to evaluate the Kambo beans, whereas participants who learned about the stereotype by reference to round beans were asked to evaluate round beans. We submitted participants' evaluations to a 3 (disconfirmation) X 2 (category label) ANOVA with both factors varying between subjects. This analysis revealed that participants evaluated the stereotype-related group more positively without ($M = 6.00$, $SEM = .162$) than with ($M = 5.15$, $SEM = .20$) the category label Kambo, $F(1, 121) = 10.95$, $p = .001$, $\eta^2 = .08$. This effect was independent of the type of disconfirmation, as suggested by the lack of an interaction of category and disconfirmation, $F < 1$. More important to the present hypotheses, there was an indication that evaluations varied as a function of disconfirmation, $F(2, 121) = 2.66$, $p = .074$, $\eta^2 = .04$ (see Figure 7). Contrast analyses suggest that the stereotype related group was evaluated more positively with clustered than with distributed disconfirmation ($p = .023$). The contrasts between no disconfirmation and clustered disconfirmation ($p = .247$), and no disconfirmation and distributed disconfirmation ($p = .276$), failed to reach statistical significance. The subtyping permitted by the clustered disconfirmation preserved the positivity of the stereotype.

Belief in Stereotype

Previous research on subtyping suggests that the opportunity to exclude disconfirming exemplars from a stereotyped group leaves the original belief in the stereotype intact. Only non-predictable disconfirmation, which prevents systematic de-categorization of exemplars, is likely

to promote a reduction of the perceived validity of the stereotype (Richards & Hewstone, 2001). To test whether this effect was also evident in the BeanFest paradigm, we analyzed participants' belief in the verbally conveyed stereotype about round beans / Kambo beans as a function of the presence of absence of a verbal category and the type of disconfirmation. The 3 (disconfirmation) X 2 (category) ANOVA with both factors varying between subjects revealed a significant main effect of category, $F(1, 121) = 4.22, p = .042, \eta^2 = .03$, indicating that participants perceived the previously conveyed stereotype to be more accurate when a verbal category was absent ($M = 6.10, SEM = .153$) than when it was present ($M = 5.55, SEM = .22$). More importantly, however, the belief was affected by the type of disconfirmation, $F(2, 121) = 3.02, p = .053, \eta^2 = .05$, see Figure 8. Contrast analyses revealed that the group who had experienced distributed disconfirmation found the verbally conveyed stereotype to be less accurate than those who experienced no disconfirmation ($p = .031$) or clustered disconfirmation ($p = .041$), with the latter two groups having comparable levels of belief ($p = .860$). This effect was independent of the presence of absence of a category, as indicated by the lack of a two-way interaction, $F < 1$.

Learning

As in Study 1, we recoded participants' responses as to whether they were correct (approaching positive beans, avoiding negative beans) or incorrect (approaching negative beans, avoiding positive beans). We submitted the overall percentage of correct responses to a 4 (Block) X 3 (Disconfirmation) X 2 (Category) ANOVA, with the first factor varying within, the last two factors varying between subjects. As in Experiment 1, we performed an arcsine transformation (Kirk, 1968) on the raw percentages before submitting them to the ANOVA. For the sake of readability, we report the raw percentages in Figure 9. The analysis revealed that participants'

performance significantly increased with the number of blocks, $F(3, 363) = 52.62, p < .001, \eta^2 = .30$. Participants improved from 82% correct responses in the first block to 91% correct responses in the fourth block. Most important to the present hypotheses, people were differentially successful depending on the type of disconfirmation, $F(2, 121) = 75.62, p < .001, \eta^2 = .56$. Contrast analyses revealed that participants performed best without disconfirmation ($M = 95\%, SEM = 1.18$), second best with clustered disconfirmation ($M = 89\%, SEM = 1.11$), and worst with distributed disconfirmation ($M = 80\%, SEM = 1.15$). This rank order was apparent in block one (all $p < .05$), block two (all $p < .001$), block three (all $p < .001$), and in block four (all $p < .01$).⁸

Performance was also better in the presence ($M = 90\%, SEM = .93$) than in the absence ($M = 86\%, SEM = .95$) of a verbal category, $F(1, 121) = 8.63, p = .004, \eta^2 = .067$. Moreover, the effect of the category differed as a function of practice, as is expressed in an interaction with the number of blocks, $F(3, 363) = 7.76, p < .001, \eta^2 = .060$. Contrast analyses indicated that the presence or absence of a category had significant effects in block one ($p < .001$) and block two ($p < .01$), but not in block three ($p = .345$) and block four ($p = .819$) (see Figure 10). Performance, however, significantly increased as a function of blocks in both groups (both $p < .001$). Apparently, the disadvantage of not having learned to discriminate round from oblong beans was overcome by practice. No other effect reached statistical significance (all $F < 3.0$, all $p > .05$). In sum, just as in Experiment 1, being confronted with clustered disconfirmation resulted in more successful behavior than being confronted with distributed disconfirmation, and this disadvantage remained significant up to the fourth block. In addition, Experiment 2 revealed that the presence of a verbal category initially improved performance, but this advantage was lost with more practice.

Discussion

The results of Experiment 2 replicated and extended what we had previously observed in Experiment 1. First, we repeated the basic experimental design, using a different stereotype (round beans instead of speckled beans) and a different set of exemplars from the bean matrix, as well as a different set of stimuli for the similarity judgments. The effects observed in Experiment 1 proved to be robust against these variations. As in the first experiment, the data were consistent with our predictions. The type of disconfirmation significantly influenced attention towards the stereotype- and disconfirmation-related dimensions. Participants gave significantly less weight to the stereotype-related dimension (i.e. shape) when disconfirmation was clustered and hence predictable than when it was distributed and hence less predictable. There was also a trend that the disconfirmation-related dimension was attended to most with clustered disconfirmation, least without disconfirmation, with the distributed disconfirmation condition falling in-between the former two.

In addition, we sought to replicate well-established findings from the literature on subtyping. While subtyping is known to maintain the belief in the stereotype as well as the central tendency of the category, disconfirmation that does not lend itself to subtyping has been repeatedly demonstrated to loosen the belief in the stereotype and to change the central tendency in the direction of the disconfirming events (Richards & Hewstone, 2001). We hypothesized that our distributed disconfirmation would be a model of the latter process, whereas clustered disconfirmation would be a model of subtyping. Consequently, we expected the evaluation of the stereotyped category (i.e. round beans) to be more positive with clustered than with distributed disconfirmation. Moreover, we expected the belief in the accuracy of the stereotype to be higher with clustered than with distributed disconfirmation. Both predictions were supported by the

results of Experiment 2. Although with both variables, the comparison between the three groups of disconfirmation only resulted in marginally significant main effects, the critical contrasts between clustered and distributed disconfirmation were significant in both cases. In addition, we examined how the willingness to apply a previously learned verbal category to the stereotyped group was affected by the experience of disconfirmation. We hypothesized that clustered and hence predictable disconfirmation would result in a tendency to exclude the disconfirming exemplars, whereas no such tendency should be apparent with distributed and hence less predictable exemplars. The data from the re-categorization task in Experiment 2 supported this prediction. What once was perfectly acceptable as a member of the Kambo category no longer was considered a Kambo following the subtyping promoted by clustered disconfirmation.

General Discussion

Encountering exemplars that contradict one's social stereotypes does not reliably promote an update of stereotypic knowledge. One psychological process that is known to prevent stereotype-change is subtyping, whereby disconfirming exemplars are excluded from the stereotypic group, keeping the perception of the group constant. While previous research has primarily studied the consequences of subtyping for the structure and stability of the original stereotype, the present research investigated how subtyping changes attention and perception. We argued that although subtyping can be driven by different motives and cognitive mechanisms, it may uniformly shape attention to and usage of perceptual dimensions related to the stereotype and subtype. Perceivers should exhibit increased attention to those features that help distinguish the regular exemplars from the disconfirming exemplars. For example, after having learned a subtype about Black businessmen, attention towards cues indicating the profession of a person such as clothing should be enhanced. This also implies, compared to the

case of no disconfirmation, a reduced attention towards those features that previously constituted the stereotype. For example, attention perceptual features related to ethnic descent such as skin color should be reduced after learning about a subtype of an ethnic group.

To test these hypotheses, we induced an artificial stereotype, relating to a certain group of stimuli in a computer game, which simulates instrumental behavior in the context of collecting positive and negative beans (Fazio et al., 2004). Importantly, the stimuli differed with respect to two simple and quantifiable perceptual dimensions, i.e. shape and the number of speckles. In two experiments, we created the equivalent of socially-transmitted stereotypes by informing participants that other students who played the game found a certain group of stimuli to be beneficial. During the game, participants either experienced no disconfirmation, clustered disconfirmation, or distributed disconfirmation of the stereotype they had acquired. The latter two conditions differed only with respect to the predictability of the disconfirmation, i.e. how easily the disconfirming exemplars could be identified by attending to the stereotype-unrelated dimension.

The results indicate that within this necessarily artificial and controlled environment, we were able to replicate a variety of basic effects documented within the subtyping literature. In particular, Experiment 2 demonstrated that subtyping was more evident with clustered disconfirmation than with distributed. In the clustered condition, the fact that a given exemplar would disconfirm was related to another salient feature such as a particular shape or degree of speckles. Participants with clustered disconfirmation maintained the belief in the stereotype and the perceived valence of the group at levels comparable to the condition without disconfirmation. Also, disconfirming exemplars were de-categorized, as indexed by our repetition of a simple category-identification task. The beans that disconfirmed the stereotype were no longer

considered to warrant the category label. Results were quite different when disconfirming exemplars were not associated with another salient feature (distributed disconfirmation). In this case, the belief in the stereotype was reduced, and the perception of the stereotypic category was more negative compared to the other two conditions. Moreover, there was no indication that the specific disconfirming exemplars were de-categorized, suggesting that they could not be recognized easily. This pattern of results suggests that the present paradigm indeed serves as a valid laboratory analogue of subtyping.

Most important, both experiments provide evidence that clustered, but not distributed disconfirmation reduced the attention towards the dimension that was related to the stereotype. The experiments also provided evidence that clustered disconfirmation actually increases the attention directed at non-stereotypic dimensions. When features other than the stereotypic facilitated making good decisions, participants started to pay more attention to them, and to have less faith in the features that were relevant according to the stereotype. With regular, but unpredictable disconfirmation, this effect was not present.

Implications for Research on Subtyping

The present results have important implications for research on subtyping and stereotype change. First, in demonstrating that subtyping increases attention to and use of subtype-related perceptual dimensions, and decreases attention to and use of stereotype-related perceptual dimensions, the present research uncovered a hitherto overlooked consequence of subtyping. This finding is important because it relates to stereotype change in various ways. On one hand, the observed attentional effects may promote maintenance of stereotypes in the face of disconfirming information. In fact, becoming more efficient at discriminating confirming and disconfirming exemplars may contribute to the preservation of the original stereotype. On the

other hand, the attentional consequences of subtyping clearly affect stereotype use, restricting its domain of application.

Recent evidence suggests that social categories may be activated not inevitably upon the perception of exemplars, but as a function of which perceptual features are attended to and, hence, how the target is categorized (Quinn & Macrae, 2005). From this perspective, one can predict that the attentional effects observed in the present research may reduce the degree to which the stereotypic category is used to categorize exemplars. Take the hypothetical example of a stereotype holder who develops a subtype of a race stereotype, i.e. Black physicians. Based on the present findings, this should enhance attention to cues that help recognizing occupation when members of the stereotyped group are encountered. This, in turn, may reduce the likelihood that racial features primarily drive categorization processes, and consequently the activation of associated stereotypic features. It is important to note that the predictability of disconfirmation may be a crucial precondition for the present attentional effects to occur. Without a consistent pairing of a particular perceptual dimension with the occurrence of stereotype-disconfirmation, giving special attention to a dimension other than the stereotype-related ones has no functional value.

The present research also highlights a source of subtyping that has received relatively little attention. Just as in the real world, many motives may have facilitated subtyping in the present study. However, the nature of the present paradigm makes it likely that subtyping was greatly influenced by the need to optimize instrumental behavior during the game. As such, subtyping in the present paradigm may be understood in terms of mechanisms of action-control in general (e.g., Hommel et al., 2001), and learning about contingencies between stimuli, behaviors, and outcomes in particular (e.g., Fiedler, Walther, Freytag, & Stryczek, 2002). This

mimics the way stereotypes are often used in natural settings: as (sometimes biased) tools that help individuals make behavioral decisions and judgments while navigating the social environment, to decide whom to trust or distrust, to hire or to fire, to socialize with to ostracize. Indeed, subtyping in the service of instrumental behavior should be particularly likely when perceivable features allow predicting whether an exemplar will or will not confirm the stereotype. For example, when interacting with a Black physician, various situational and personal features (e.g., consulting the person at a practice; the person wearing a white coat) indicate that the person belongs to the hypothetical Black physician subtype, thereby presumably triggering a different set of behavioral scripts than when interacting with a non-subtyped Black person. Without predictability, subtyping would not affect action control in instrumental settings. Interestingly, subtyping in the service of other motives, such as maintaining a societal status quo or protecting existing category structures from change may not be as dependent on predictability. Potentially, these two forms of subtyping differ also with respect to their effects on attention. Future research should aim at comparing subtyping in the service of action control and subtyping in the service of other motives.

Validity of the Present Paradigm

The present research differs from most previous studies on subtyping in at least three ways. First, we used artificial instead of naturalistic stereotypes and stimuli. Second, the nature of the learning task required, in contrast to many other studies, the exclusive use of a positive stereotype. Third and relatedly, our participants formed verbally conveyed preconceptions about rewards that were associated with Kambo beans, and disconfirming exemplars challenged the group-outcome association. In previous studies, however, participants had preconceptions about traits (e.g., introversion) that could be associated with certain groups (e.g. librarians), and the

disconfirming exemplars challenged the group-trait association. Given these differences, one may wonder about the extent to which the present results may generalize to other settings.

For a number of reasons, we are quite confident that such generalizations would be valid. First, the findings clearly indicate that we were able to replicate established empirical markers of subtyping in the present paradigm. Clustered compared to distributed disconfirmation led to greater belief in the stereotype, a more extreme central tendency, and a de-categorization of disconfirming exemplars. Thus, similar processes apparently operated in this artificial task as in more naturalistic paradigms. Second, just as with stereotyping in general (e.g., Smith & DeCoster, 1998), subtyping can be partially understood as a consequence of general principles of learning and memory (Queller & Smith, 2002). Particularly, Queller and Smith (2002) demonstrated that subtyping is an inherent feature of connectionist network models of memory. In their simulations, networks were fed with inputs designed to mimic stereotypes and stereotype disconfirmation. But the inputs are abstract patterns of activation values ranging between -1 and +1. As such, they may represent any category structure, be it social or non-social. Hence, connectionist learning-mechanisms exhibit subtyping independent of which specific category or stereotype is actually represented. In line with this reasoning, artificial stimuli have been previously used to successfully study phenomena related to stereotyping or subtyping (e.g., Rothbart & Lewis, 1988). Third, and related to the previous argument, only a few of the potential motives for subtyping are specific to social stimuli. In fact, such specificity is likely only for such motives as the desire to maintain social stereotypes justifying inequalities or the desire to act in a fair manner. The motive to think and act consistently, the need for closure, needs to explain the disconfirmation of expectancies, and the need to discriminate between confirming and disconfirming exemplars so as to optimize instrumental behavior, on the other hand, are

neither theoretically, nor empirically limited to social settings and stimuli (for examples, see Kunda & Oleson, 1995; Shepard, 1987).

Moreover, we wish to reiterate the point we made in the introduction section about positive versus negative stereotypes. When negative prejudices operate, people are relatively less likely to experience disconfirmation than when positive prejudices exist (Fazio et al., 2004), because they presumably try to minimize contact with the negatively-valued group. However, when interactions with a negatively-stereotyped group do occur, any ensuing instances of disconfirmation will promote the very same attentional shift processes that have been illustrated in the present research. Hence, the fact that disconfirmation is less likely given the avoidance behavior prompted by negative attitudes does not mean that the present findings are irrelevant to negative stereotypes. The relative infrequency of such disconfirming events simply suggests a more prolonged process. Thus, the intensity of the effect, not its occurrence, may be less for negative stereotypes than for positive stereotypes.

Finally, introducing positive and negative outcomes instead of stereotypic traits is an alteration that should have little influence on the processes underlying subtyping and its consequences. Indeed, we observed substantial convergence between our findings and the conditions known to induce subtyping in traditional paradigms. One reason for this convergence could be the structural similarities between our set-up and previous studies. In traditional studies, a stereotyped group label (e.g., Black) is associated with defining characteristics (e.g., dark skin tone) and stereotypic traits (e.g., athletic). Disconfirming individuals (e.g., a non-athletic Black person) challenge the link between the group label (e.g., Black) and the stereotypic trait (e.g., athletic), and the disconfirming individual may be decategorized. The link between defining characteristics (e.g., dark skin tone) and the group label (e.g., Black), however, is not challenged

explicitly. In our set-up, a stereotyped group label (Kambo beans) is associated with a defining characteristic (e.g., round shape) and particular outcomes (reward when approached).

Disconfirming exemplars (a non-rewarding Kambo bean) challenge the link between the group label (e.g., Kambo) and the expected outcome (reward). As in the traditional case, the disconfirming exemplar may be decategorized, while the link between defining characteristic (e.g., round shape) and the group label (Kambo beans) is not challenged explicitly. Given these similarities, we think it is obvious that the two paradigms result in similar outcomes. If anything, we would expect effects to be stronger with a valence-based disconfirmation, given that valence generally has strong effects on attention and categorization (e.g., Fazio & Dunton, 1997; Smith, Fazio, & Cejka, 1996). In sum, there are many reasons to believe that the present findings reflect fundamental processes that are applicable to more naturalistic settings involving more social stimuli.

Conclusion

In essence, the present research demonstrates a hitherto neglected cognitive consequence of subtyping: Subtyping enhances the attention to and use of perceptual dimensions related to the subtype, and reduces the attention to and use of perceptual dimensions related to the stereotype. This finding is relevant for at least two reasons. First, it implies that subtyping, besides stabilizing the original stereotype, may have the potential to reduce the use of the original category associated with the stereotype. Second, it provides further insights into the mechanisms driving subtyping. Attentional effects of the kind we observed in the present studies may be important cognitive mediators that facilitate the construction of subtypes. Moreover, the use of a learning-paradigm makes it likely that subtyping was partially driven by the need to optimize

instrumental behavior. This motive, in parallel with other social and non-social motives, may also be relevant in naturalistic settings, where stereotype holders may use subtype-membership as a discriminant cue for selecting social behavior adjusted to yield maximally rewarding interactions with the subtyped exemplars. Finally, the present research suggests that the present learning-paradigm may serve as a laboratory model of subtyping. Despite its artificial nature, we were able to replicate central empirical markers of subtyping, suggesting that comparable psychological mechanisms were in operation.

In their review of the literature on subtyping, Richards and Hewstone (2001) concluded: “For subtyping, the consequences are largely negative because ... this process leaves the content of the stereotype unchanged and excludes disconfirming members from the group” (p. 70). The present research extends this assessment. The findings illustrate another cognitive consequence, one which is of less dubious value. Subtyping decreases attention to stereotype-related dimensions, increases attention to other dimensions, and may therefore reduce the use of the original stereotypic category.

References

- Brewer, M. B., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology*, 41, 656-670.
- Cumming, G., & Finch, S. (2005). Inference by eye: confidence intervals and how to read pictures of data. *American Psychologist*, 60, 170-80
- Dovidio, J. F., Gaertner, S. L., & Kawakami, K. (2003). Intergroup contact: The past, present, and future. *Group Processes and Intergroup Relations*, 6, 5-21.
- Fazio, R. H., & Dunton, B. C. (1997). Categorization by race: The impact of automatic and controlled components of racial prejudice. *Journal of Experimental Social Psychology*, 33, 451-470.
- Fazio, R. H., Eiser, J. R., & Shook, N. J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology*, 87, 293-311
- Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row Peterson.
- Fiedler, K., Walther, E., Freytag, P., & Stryczek, E. (2002). Playing mating games in foreign cultures: A conceptual framework and an experimental paradigm for trivariate statistical inference. *Journal of Experimental Social Psychology*, 38, 14-30.
- Hewstone, M., Macrae, C. N., Griffiths, R., Milne, A., & Brown, R. (1994). Cognitive models of stereotype change: (5) Measurement, development, and consequences of subtyping. *Journal of Experimental Social Psychology*, 30, 505-526.
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The theory of event coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24, 849-878.

Judd, C. M., & Park, B. (1993). Definition and assessment of accuracy in social stereotypes.

Psychological Review, 100, 109-128.

Kirk, R. E. (1968). *Experimental design: Procedures for the behavioral sciences*. Belmont, CA:

Brooks/Cole.

Kruglanski, A. W., & Freund, T. (1983). The freezing and un-freezing of lay inferences: Effects on impression primacy, ethnic stereotyping and numerical anchoring. *Journal of*

Experimental Social Psychology, 19, 448-468.

Kunda, Z., & Oleson, K. C. (1995). Maintaining stereotypes in the face of disconfirmation:

Constructing grounds for subtyping deviants. *Journal of Personality and Social Psychology*, 68, 565-579.

Kunda, Z., & Oleson, K. C. (1997). When exceptions prove the rule: How extremity of deviance determines the impact of deviant exemplars on stereotypes. *Journal of Personality and*

Social Psychology, 72, 965-979.

Levine, G. M., Halberstadt, J. B., & Goldstone, R. L. (1996). Reasoning and the weighting of attributes in attitude judgments. *Journal of Personality and Social Psychology*, 70, 230-240.

Lott, B. (2002). Cognitive and behavioral distancing from the poor. *American Psychologist*, 57, 100-110.

Maurer, K. L., Park, B., & Rothbart, M. (1995). Subtyping versus subgrouping processes in stereotype representation. *Journal of Personality and Social Psychology*, 69, 812-824.

Nosofsky, R. M. (1992). Similarity scaling and cognitive process models. *Annual Review of Psychology*, 43, 25-53.

- Olson, M. A., & Fazio, R. H. (2004). Trait inferences as a function of automatically-activated racial attitudes and motivation to control prejudiced reactions. *Basic and Applied Social Psychology, 26*, 1-11.
- Queller, S., & Smith, E. R. (2002). Subtyping versus bookkeeping in stereotype learning and change: Connectionist simulations and empirical findings. *Journal of Personality and Social Psychology, 82*, 300-313
- Quinn, K. A., & Macrae, C. N. (2005). Categorizing others: The dynamics of person construal. *Journal of Personality and Social Psychology, 88*, 467-479.
- Richards, Z., & Hewstone, M. (2001). Subtyping and subgrouping: Processes for the prevention and promotion of stereotype change. *Personality and Social Psychology Review, 5*, 52-73.
- Rothbart, M., & Lewis, S. (1988). Inferring category attributes from exemplar attributes: Geometric shapes and social categories. *Journal of Personality and Social Psychology, 55*, 861-872.
- Shepard, R. N. (1987). Toward a universal law of generalization for psychological science. *Science, 237*, 17-23.
- Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology, 74*, 21-35.
- Smith, E. R., Fazio, R. H., & Cejka, M. A. (1996). Accessible attitudes influence categorization of multiply categorizable objects. *Journal of Personality and Social Psychology, 71*, 888-898.

- Taylor, S. E. (1981). A categorization approach to stereotyping. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 83–114). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Towles-Schwen, T., & Fazio, R. H. (2003). Choosing social situations: The relation between automatically-activated racial attitudes and anticipated comfort interacting with African Americans. *Personality and Social Psychology Bulletin*, 29, 170-182.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961–977.
- Yzerbyt, V., Coull, A., & Rocher, S. (1999). Fencing off the deviant: The role of cognitive resources in the maintenance of stereotypes. *Journal of Personality and Social Psychology*, 77, 449-462.

Author Note

This research was supported by a grant from the German Science Foundation (DFG) to Roland Deutsch (De 1150/1-1) and a grant from the National Institute of Mental Health to Russell H. Fazio (MH38832). We thank the Social Cognition Research Group at the Ohio State University, the members of the Person Memory Interest Group of 2006, and Robert Kordts for their thoughtful comments on this research. Correspondence concerning this article should be addressed to Roland Deutsch, who is now at the University of Würzburg, Lehrstuhl für Psychologie II, Röntgenring 10, 97070 Würzburg, Germany. Email: deutsch@psychologie.uni-wuerzburg.de.

Footnotes

1. Besides subtyping, which is the focus of the present work, stereotype-disconfirmation can also trigger a process called subgrouping. In this case stereotype-disconfirmation is processed in a way that promotes the formation of subgroups, and can increase the perceived variability of the category and change its central tendency if it. Unlike subtyping, subgrouping is not limited to individuals who disconfirm the stereotype, and the subgroups are not excluded from the original category. Instead, “subgroups may as likely be formed for clusters of individuals who are perfectly consistent with the group stereotype but who manifest the stereotype in some unique and different way” (Maurer et al., 1995; p. 813). Although superficially similar to subtyping, subgrouping occurs under different conditions and has quite different cognitive consequences (for a review, see Richards & Hewstone, 2001).
2. For applications of MDS within the social psychological literature, see Levine, Halberstadt, and Goldstone (1996), who employed the technique to examine the impact of analyzing reasons for one’s liking on the weighting of stimulus dimensions, and Fazio and Dunton (1997), who used MDS to study categorization by race as a function of automatically-activated racial attitudes.
3. When comparing two means that stem from independent samples, the overlap of 95% CIs provides an estimate of whether the two means differ significantly. According to the “rules of eye” delineated by Cumming and Finch (2005), two independently sampled means differ at $p \leq .05$ if the two error bars do not overlap more than half of their length. For further information on the interpretation of CIs and error bars, please see Cumming and Finch (2005).
4. Recall that the pretesting we had mentioned earlier indicated the naturally dominant dimension to be shape. That is, when participants who did not play BeanFest and, hence, were not exposed

to any differential outcomes for the various beans, judged similarity, they attended much more to the shapes dimension than to the speckles dimension. Thus, the relative equivalence of the weights for shape and speckles in the no disconfirmation and distributed disconfirmation conditions appears to reflect greater attention than is typical for speckles when that dimension has been made more hedonically significant. In the actual game, speckles distinguished those beans that produced gains from those that yielded losses.

5. We also analyzed the squared dimension weights including the third dimension that was related to novelty, using a 3 (dimension) X 3 (disconfirmation) ANOVA. The main effect of dimension was significant, $F(2, 132) = 91.124, p < .001, \eta^2 = .58$, indicating that the least weight was given to the novelty dimension. This main effect was qualified by an interaction of dimension and disconfirmation, $F(4, 132) = 4.23, p = .003, \eta^2 = .12$. Contrast analyses revealed that the novelty dimension was not affected by disconfirmation, $F(2, 66) = 2.31, p = .108, \eta^2 = .065$, whereas the shape dimension, $F(2, 66) = 3.78, p = .028, \eta^2 = .103$, and the speckles dimension, $F(2, 66) = 4.00, p = .023, \eta^2 = .108$, were. The main effect of disconfirmation was not significant, $F < 1$. This further suggests that the novelty dimension is relatively unimportant.

6. There was also evidence that the linear effect of the number of blocks was different for the three types of disconfirmation, $F(6, 198) = 2.71, p = .015, \eta^2 = .08$. Contrast analyses indicate that for the no disconfirmation group, performance did not differ as a function of block, reflecting the close to perfect performance from the first block on, $F(3, 64) = 2.15, p = .102, \eta^2 = .092$. For the clustered disconfirmation group, performance varied as a function of blocks, $F(3, 64) = 13.81, p < .001, \eta^2 = .39$. Simple contrasts indicate that performance increased between block one and two ($p < .01$), as well as block three and four ($p < .05$), but not between block two and three ($p = .126$). For the distributed disconfirmation group, performance also varied as a

function of blocks, $F(3, 64) = 6.36, p = .001, \eta^2 = .23$. Simple contrasts indicate that with distributed disconfirmation, performance increased between block one and two ($p < .001$), but not between block two and three ($p = .775$) or between between block three and four ($p = .216$).

7. We also analyzed the squared dimension weights including the third dimension that was related to novelty, using a 3 (dimension) X 3 (disconfirmation) X 2 (category label present or absent) ANOVA. The main effect of dimension was significant, $F(2, 242) = 151.04, p < .001, \eta^2 = .55$, indicating that the least weight was given to speckles, the most weight was given to shape, and the novelty dimension fell in-between. This main effect was qualified by an interaction of dimension and disconfirmation, $F(4, 242) = 3.08, p = .017, \eta^2 = .048$. Contrast analyses revealed that the novelty dimension was not affected by disconfirmation, $F(2, 121) = 1.09, p = .339, \eta^2 = .018$, whereas the shape dimension, $F(2, 121) = 3.89, p = .023, \eta^2 = .06$, and the speckles dimension, $F(2, 121) = 2.73, p = .066, \eta^2 = .044$, were. No other effect was significant, all $F < 1.3$. As in Experiment 1, this suggests that the novelty dimension is relatively unimportant.

8. There was also evidence that the linear effect of the number of blocks was different for the three types of disconfirmation, $F(2, 121) = 4.02, p = .001, \eta^2 = .062$. Further analyses indicate that the effect of block was significant in all three groups (all $p < .01$), but that the effect of block was strongest with clustered disconfirmation ($\eta^2 = .460$), weakest without disconfirmation ($\eta^2 = .117$), and intermediate for distributed disconfirmation ($\eta^2 = .259$). For the clustered disconfirmation group, simple contrasts indicate that performance increased between block one and two ($p < .001$), as well as between block two and three ($p < .001$), but not between block three and four ($p = .250$). Without disconfirmation, simple contrasts indicate that performance increased between block one and two ($p < .01$), between block two and three ($p = .01$), and between block three and four ($p = .038$). For the distributed disconfirmation group, simple

contrasts indicate that performance increased marginally significant between block one and two ($p < .064$), and marginally significant between block three and four ($p = .083$), but not between block two and three ($p = .532$).

Table 1

Beta weights from regression analyses predicting each of the three MDS dimensions from the beans' shape, speckles, and novelty (Experiment 1).

Dependent Variables	Predictors		
	Shape	Speckles	Novelty
Dimension 1	-0.839	0.266	0.026
Dimension 2	0.115	-0.966	-0.324
Dimension 3	-0.052	0.205	0.866

Note: Beta weights in boldface are significant at at least $p < .01$.

Table 2

Beta weights from regression analyses predicting each of the three MDS dimensions from the beans' shape, speckles, and novelty (Experiment 2).

Dependent Variables	Predictors		
	Shape	Speckles	Novelty
Dimension 1	-0.971	-0.015	-0.209
Dimension 2	-0.139	-0.016	-0.982
Dimension 3	0.131	-0.942	0.039

Note: Beta weights in boldface are significant at at least $p < .01$.

Figure Captions

Figure 1. Overview over the 10 by 10 bean matrix used as a basis for generating stimulus materials in Experiments 1 and 2.

Figure 2. Example of a simple reward-structure (A), no disconfirmation within the positive area (B), clustered disconfirmation within the positive area (C), and distributed disconfirmation within the positive area (D).

Figure 3. Squared individual weights of the dimensions *shape* and *speckles* derived from an MDS in Experiment 1. *Speckles* was the dimension related to the positive stereotype. Error bars reflect 95% confidence intervals (CIs)³ of the means.

Figure 4. Percentage of correct responses (avoiding negative and approaching positive beans) towards the beans (Experiment 1). Error bars reflect 95% CIs of the means.

Figure 5. Squared individual weights of the dimensions *shape* and *speckles* derived from an MDS in Experiment 2. *Shape* was the dimension related to the positive stereotype. Error bars reflect 95% CIs of the means.

Figure 6. Probability of applying the stereotype-related category (*Kambo*) to stereotype confirming vs. stereotype disconfirming exemplars. Error bars reflect 95% CIs of the means.

Figure 7. Evaluation of Kambo beans / round beans as a whole as a function of disconfirmation. Error bars reflect 95% CIs of the means.

Figure 8. Judgments of perceived stereotype-validity as a function of disconfirmation. Error bars reflect 95% CIs of the means.

Figure 9. Percentage of correct responses (avoiding negative and approaching positive beans) towards the beans as a function of the level of disconfirmation (Experiment 2). Error bars reflect 95% CIs of the means.

Figure 10. Percentage of correct responses (avoiding negative and approaching positive beans) towards the beans as a function of the presence or absence of a verbal category (Experiment 2). Error bars reflect 95% CIs of the means.

Fig.1




Fig.2

-8	-7									6
-10									9	
									10	8
	-7									7
-8	-10								10	
-10									8	10
	-7									7
-10	-8									
-9									10	
	-6								8	7

	6
9	
10	8
	7
10	
8	10
	7
10	
8	7

	6
9	
10	8
	7
10	
8	10
	7
-10	
-8	-7

	6
9	
10	-8
	7
-10	
8	10
	7
10	
8	-7

Fig.3




Fig.4




Fig.5




Fig.6




Fig.7




Fig.8




Fig.9




Fig.10

