Journal of Personality and Social Psychology
1989, Vol. 56, No. 6, 866-875

0022-3514/89/$00.75

Category Learning and Change: Differences in Sensitivity to Information That Enhances or Reduces Intercategory Distinctions

Joachim Krueger, Myron Rothbart, and N. Sriram
University of Oregon

We examined the amenability of abstractions of categories to new and relevant information. In Experiment 1, Ss formed impressions of 2 sets of numbers by periodically estimating the cumulative means of each set. During the 1st half of the procedure, the 2 means were mathematically stable. During the 2nd half of the procedure, the mean of 1 set was modified and the mean of the other set remained unchanged. We predicted and found that the resultant estimates for the modified category changed more when the mean difference between the 2 categories was enhanced than when it was reduced. Experiment 2 suggested that the accentuation effect results from a 2-stage process of category learning (Stage 1) and category change (Stage 2). Experiment 3 replicated the effect with person categories. The relevance of category accentuation is discussed with respect to the modifiability of social beliefs.

> The mind must think with the aid of categories. Once formed, categories are the basis for normal prejudgment. We cannot possibly avoid this process. Orderly living depends on it. (Allport, 1954, p. 19).

One of the enduring problems of social psychology concerns the nature of social beliefs. How are beliefs formed and how are they modified by experience? When do beliefs mirror reality with sufficient accuracy and when do they deviate from it? More than 30 years ago, Gordon Allport, in his classic book on prejudice, maintained that processes of categorization play an important role in the organization and simplification of social information (Allport, 1954). Categorization provides a parsimonious representation of knowledge by aiding "the search for meaning, the need to understand, [and] the trend toward better organization of perceptions and beliefs" (Katz, 1960). Can the attainment of structure be accomplished without any costs? Under which conditions are categories veridical representations of reality, and are they sufficiently responsive to new and relevant information?

In this article, we review some of the ways in which categorization processes may affect the formation of social beliefs. Then we introduce an experimental paradigm that permits the assessment of the degree to which new information, which modifies the attributes of a category, is reflected in perceptions of the attributes of that category.

The studies reported in this article constitute, in part, the dissertation research of Joachim Krueger. Completion of this research would not have been possible without the support from Grant MH40662-02 from the National Institutes of Mental Health to Myron Rothbart.

We are indebted to Asher Cohen, Morton Gernsbacher, Lewis Goldberg, Douglas Hintzman, Oliver John, and Steven Keele for their comments and encouragement.

Correspondence concerning this manuscript should be addressed to Joachim Krueger, who is now at the Max Planck Institute for Human Development and Education, Lentzeallee 94, D-1000 Berlin 33, Federal Republic of Germany.

Contrast and Accentuation Effects

Numerous studies in cognitive and social psychology have shown that categorization can lead to systematic perceptual and judgmental biases. Dawes, Singer, and Lemons (1972), for example, showed that real intergroup differences are exaggerated when the opposing categories are salient. These authors had subjects who were self-defined as hawks and doves regarding the Vietnam War construct statements that they believed hawks and doves would endorse. When rating hawkish and dovish attitude statements on political extremity, statements attributed to the in-group were perceived as less extreme than statements attributed to the out-group. Moreover, when in-group members were asked to write statements they would consider to be typical for the out-group, they generated items that were rejected by the out-group as being too extreme.

Dawes et al. (1972) interpreted these findings as evidence for a perceptual contrast effect. Using the attitudes of their own group as an anchor, subjects displaced the attitudes of the out-group toward the opposite end of the political spectrum. Studies in the minimal-group paradigm have shown that different perceptions of groups may arise even without real a priori group differences (Allen & Wilder, 1979; Billig & Tajfel, 1973; Rabbie & Horwitz, 1969). Allen and Wilder (1979) randomly categorized subjects as admirers of Klee paintings and admirers of Kandinsky paintings and found that different perceptions of the average attributes of the groups emerged. Between-category boundaries were exaggerated by ascribing different characteristics to groups that bore different labels. Specifically, more positive attributes were ascribed to in-groups than to out-groups. In sum, these studies demonstrated the propensity to perceive intergroup contrasts that favor the in-group.

Tajfel and Wilkes (1963), in a seminal article, argued that the exaggeration of intercategory differences is a general perceptual phenomenon that is limited neither to differences in category favorability, nor to the involvement of the judge in one of the groups. Tajfel and Wilkes (1963) proposed a top-down perceptual process, in which the perception of each stimulus is affected by the knowledge of the category membership of this stimulus. Subjects judged the length of eight graded lines. Each line was labeled either *A* or *B*. In one condition, the categorical label was assigned randomly. In another condition, the shorter four lines shared one label and the longer four shared the other; that is, a second, dichotomous variable was superimposed on the continuous variable of length. In this condition, subjects exaggerated the length of the shorter lines in the long category and minimized the length of the longer lines in the short category; that is, intercategory boundaries were accentuated. Similarly, Eiser (1971) showed that the evaluation of attitudinal statements was polarized when a categorical variable (source of statement) was superimposed on a continuous, evaluative dimension. Using statements of varying degrees of permissiveness toward drug use as the continuous variable and two newspaper sources as the categorical variable, Eiser found that permissive statements were judged as more permissive if they were all attributed to the same newspaper, and that nonpermissive statements were judged as more punitive if they were all attributed to the same newspaper.

The Tajfel and Wilkes (1963) and Eiser (1971) research demonstrated the effects of categorization on the perception of individual exemplars of categories. It did not explore the perception of category attributes. Using the framework of social-judgment theory (Hovland, Harvey, & Sherif, 1957), Wilder and Thompson (1988) examined assimilation and contrast effects on the evaluations of out-groups. In a mock-jury paradigm, members of a fictitious moderate out-group rendered a verdict that was different from that of the in-group (the subjects). The degree of this difference was varied across three levels. Half the subjects were also exposed to a second fictitious out-group that espoused an extremely discrepant position. The results showed that the perceived positions of the in-group and the extreme out-group served as anchors for the liking of the moderate out-group. When the moderate out-group's position was more similar to that of the in-group than to that of the extreme out-group, it was assimilated into the in-group and liked more. When its position was more similar to the extreme out-group, it was contrasted away from the in-group and liked less.

Wilder and Thompson's (1988) research elucidated context effects on the perception of between-category differences, but like most studies in this area, it used a static approach that did not address ongoing category learning. Do accentuation effects occur as new elements are incorporated into existing categories? To gauge the effects of categorization on changes of judgment, a dynamic procedure is needed that permits the sequential presentation of category examplars and the trial-by-trial assessment of category representation (Busemeyer & Myung, 1988).

Paradigm for Category Change

In the present research, we introduce an experimental method for monitoring category learning over time. We hypothesize that when people have learned to distinguish between two categories, they will later incorporate new information in a way that enhances rather than reduces perceived intercategory distinctions. As a thought experiment, consider a person who has learned that, on the average, members of two groups (A and B) differ in the degree of their sociability, with Bs being more sociable than As. Now imagine two possible conditions. In the first condition, an observer encounters new members of Group A who are more sociable than most As. This information should increase the perceived average sociability of Group A by some quantity, Q, making Group A more similar to Group B. In an alternative condition, an observer encounters new members of Group A who are less sociable than most As. This information should reduce the perceived average sociability of A by the same quantity, Q, thus making Group A less similar to Group B. In principle, the absolute magnitude of change in Group A should be the same, but if intercategory differences are accentuated perceived change should be greater in the second condition than in the first.

The research we report in this article reflects the start of a series of experiments that deal with category learning and change. Complex social beliefs are characterized by a number of different dimensions, and our long-term goal in this research is to assess the contribution of these dimensions. For example, it is possible to give information about unfamiliar categories that does not elicit an affective response. One can also learn new information about familiar categories, in which strong expectations and affective orientations are involved. The strategy we have adopted is to begin with simple beliefs, where subjects learn abstract properties about novel and affectively neutral categories, and later to increase complexity by changing the properties from abstract to more concrete attributes and by using familiar rather than novel categories.

Our paradigm allows continuous monitoring of the degree to which subjects incorporate new information that modifies the attributes of a category. Our hypothesis is that when the new elements enhance the average intercategory difference, judgments of the category mean will change more than when the new elements reduce the intercategory difference.

In all three experiments, we used three-digit numbers as stimuli because they offer several advantages. First, numbers are unambiguous. For example, when the stimulus is 173, it is unlikely to be misperceived as a different number. That is, if accentuation effects are obtained on the level of category means, they can only be attributed to a biased averaging process. Unambiguous stimuli preclude the assimilation of exemplars to expectations associated with the category. Category accentuation, therefore, cannot arise from the accurate averaging of misperceived numbers. Second, numbers permit an estimation of the mean as a measure of category representation. Studies on statistical intuition have documented a high level of accuracy in mean estimation (Beach & Swensson, 1966; Spencer, 1963). Malmi and Samson (1983) found mean estimates to be accurate even when two series of numbers were presented at the same time and each individual stimulus was shown only for 0.5 s. Third, the number continuum does not connote different degrees of favorability, and there is no category of numbers to which a subject might develop a sense of in-group loyalty. Hence, category learning will not be confounded with biases of in-group favoritism. In sum, the intuitive statistics paradigm presents a uniquely conservative test of the category-accentuation hypothesis.

In Experiment 3, the numbers represent manifestations of a person variable associated with two groups of athletes. On the

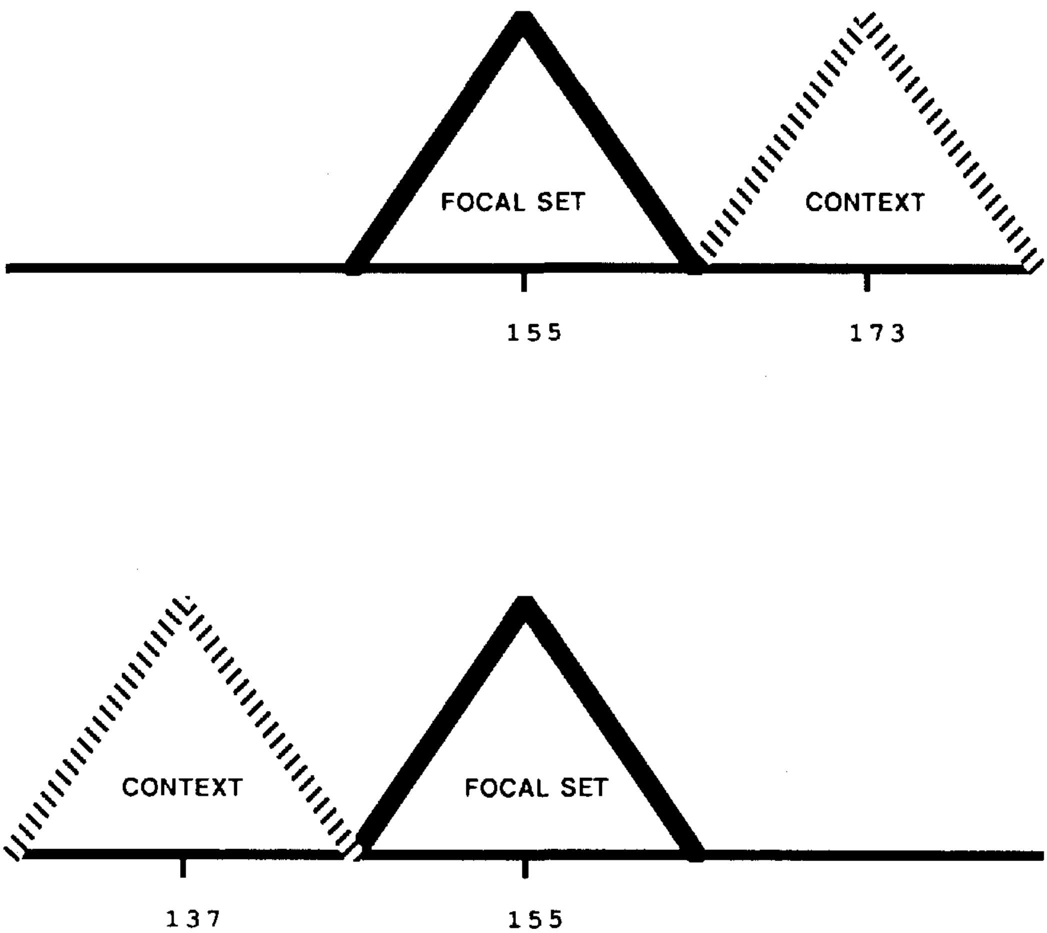

Figure 1. Schematic representation of the number distributions of the focal and contextual categories after Phase 1: Experiment 1.

basis of the presented exemplars, subjects learn a difference between the average body weights of two categories of runners that is congruent with a prior expectation.

Experiment 1

Overview

In the first phase of a two-phase learning experiment, we presented subjects with a sequence of 48 three-digit numbers, half of which were associated with one category and half with another. The numbers in each category were distributed unimodally. The two distributions did not overlap but bordered one another, and their variances were equal. After each number was presented, the subject typed that number into a computer, along with a letter identifying its category. After each set of eight trials, the subject estimated the cumulative average for each of the two categories.

In the second phase, subjects received an additional 48 numbers, 24 in each category. For one category (hereafter called the *focal category*), the second set of numbers either increased or decreased the central tendency of the original set. For the other category (hereafter called the *contextual category*), the second set of numbers was identical to the first set. The subjects' task in the second phase was identical to their task in the first phase: They entered each number into the computer, along with its category label, and they estimated the average for each category on the basis of all the numbers presented up to that point, including those presented during the category-learning phase.

There were four between-subjects experimental conditions. For half the subjects, the information in the category-change phase increased the central tendency of the focal category, and for half the subjects the information in the category-change phase decreased this central tendency. Within each of these conditions, the mean of the contextual category was either larger or smaller than the mean of the focal category. Thus, it was possible to compare subjects' perceptions of increasing and decreasing changes of the mean of a distribution when those changes reduced or enhanced the average difference between the two categories.

The category-accentuation effect, if obtained, may result from two distinct processes. One possibility is that subjects resist the loss of a clear distinction by ignoring new information that reduces intercategory differences. That is, perceived mean changes might be inhibited when the contextual category is located in the direction of change. Alternatively, subjects may be particularly attentive to new information when it enhances the difference between two categories, that is, when the mean in the focal category moves away from the contextual category. We constructed a baseline condition in which the mean of the contextual category was far removed from the focal category. Thus, the degree of context-free perceived category change could be assessed and compared with the two experimental conditions. If the change in the baseline condition was comparable to the change in the enhancement condition but greater than the change in the reduction condition, change was inhibited in the reduction condition. If the change in the baseline condition was equal to the change in the reduction condition and both were smaller than the change in the enhancement condition, that change was facilitated in the enhancement condition.

Method

Subjects. Eighty undergraduate students enrolled in introductory psychology classes at the University of Oregon participated in this study in exchange for extra credit in their course.

Procedures. Subjects participated in groups of 4, and the sessions lasted for 50 min. The experimenter explained that the study was concerned with the intuitive estimation of averages of sequences of numbers. He pointed out that people can make sound intuitive judgments about averages without making explicit computations, and he assured the subjects that the experiment was not a test of numerical ability. Subjects worked individually on personal computers that were programmed to present 96 three-digit numbers for 2 s each. Each number was an exemplar of one of two categories. Category membership was defined by the type font in which the number was presented.

After each presentation, subjects typed in the number and entered a letter that identified the type font. This procedure ensured continuing attention to the two critical features of the stimuli: the value of the number and the type font. The 96 numbers comprised 12 blocks of 8 numbers. Within each block there were 4 numbers of each category, presented in random order. After each block, subjects estimated the mean separately for each category. Each time they made this judgment, subjects were instructed to consider all the numbers they had seen up to that point. Hence, the task became increasingly difficult as the experiment proceeded: After the 1st block, a total of 4 numbers had been presented in each set; after the 2nd block, the total was 8 numbers in each set; and by the end of the experiment, each set had 48 numbers. After the 12th block, subjects were thanked and debriefed.

Design and stimulus materials. We refer to the first half of the experiment (i.e., Blocks 1–6) as the *category-learning phase.* By the end of Block 6, 24 numbers had been presented in each category. The two distributions of numbers were constructed to have a mean difference of 18 points. Both distributions were unimodal, symmetrical, and had a standard deviation of 3.83. The two distributions bordered each other but did not overlap. Figure 1 shows a schematic representation of the two distributions after Phase 1 for the condition with large contextual numbers (top) and for the condition with small contextual numbers (bottom).

The second half of the experiment (Blocks 7–12) constituted the category-change phase. In this phase, the same numbers used in the category-learning phase were repeated in a newly randomized order for the contextual category. Thus, the cumulative mean of this category remained stable across the two phases of the experiment. For the focal category, only 8 of the 24 numbers were repetitions of earlier presentations. Randomly intermixed with these 8 numbers were 16 other numbers, which were on the average either 18 points higher or lower than the mean at the end of Phase 1. By the end of Phase 2, the grand mean of the focal category had changed by 6 points from the mean at the end of Phase 1. Subjects were not informed about the distinction between the category-learning phase and the category-change phase.

For all subjects, the mean of the focal category was 155.5 after Phase 1. For half the subjects, the contextual category was composed of smaller numbers (M = 137.5); for the other half, the contextual category was larger (M = 173.5). We also varied the direction of the mean change of the focal category so that the mean of the focal category at the end of Phase 2 was higher (M = 161.5) for half the subjects and lower (M = 149.5) for the other half. When the contextual mean was greater than the focal mean, an upward shift resulted in a reduction of the intercategory difference; a downward shift resulted in an enhancement of the intercategory difference. Conversely, when the contextual mean was smaller than the focal mean, an upward shift produced an enhancement (i.e., the distance between the two means increased) and a downward shift produced a reduction (the distance between the two means diminished). Figure 2 shows a schematic representation of the two distributions after Phase 2 for the condition with large contextual numbers (top) and for the condition with small contextual numbers (bottom).

We constructed a baseline condition to assess the perceived mean change when the contextual category (M = 805.5) was far removed from the focal category. The characteristics of the focal category were identical to those used in the experimental conditions. As in the two experimental conditions, for half the baseline subjects the mean of the focal category was higher at the end of Phase 2 than at the end of Phase 1; for the other half it was lower. With the extreme difference between the focal and the contextual means it seemed unlikely that the contextual category would affect judgments in the focal category. In sum, there were three variables in the design: Direction of change of focal category (up vs. down), condition (reduction vs. enhancement vs. baseline), and phase (learning vs. change).

Dependent variables. The first dependent variable was the estimated means of the focal category, assessed at 12 different times throughout the experiment. For analyses, we averaged the estimated means within Phase 1 and the means within Phase 2. As a second dependent variable, we subtracted the averaged means in Phase 2 from the average of the means in Phase 1. The absolute value of this difference score indicated the magnitude of average change, independent of the direction on the numerical scale.

Results

The first 3 estimates in Phase 1 were discarded as practice trials. We averaged the estimates of Blocks 4, 5, and 6 for each subject to provide a reliable measure of performance in the category-learning phase. The true means were averaged in the same fashion to provide a criterion for accuracy. We averaged the estimates of Blocks 7–12 to provide the same measure for the category-change phase. Table 1 shows averages and standard deviations for the six experimental conditions.

At the end of the category-learning phase, the averaged estimates of the mean of the focal category did not differ significantly from the true mean. Furthermore, mean estimates in the experimental conditions were not displaced away from the contextual category. Estimates in the presence of a low contextual category (M = 154.18) were not greater than estimates in the presence of a high contextual category (M = 155.67).

To test the hypothesis of category accentuation, we performed a two-way analysis of covariance (ANCOVA) on the average estimated means in the focal category in Phase 2, with the averaged estimates in Phase 1 as the covariate. The between-subjects variables were direction of change (up vs. down) and

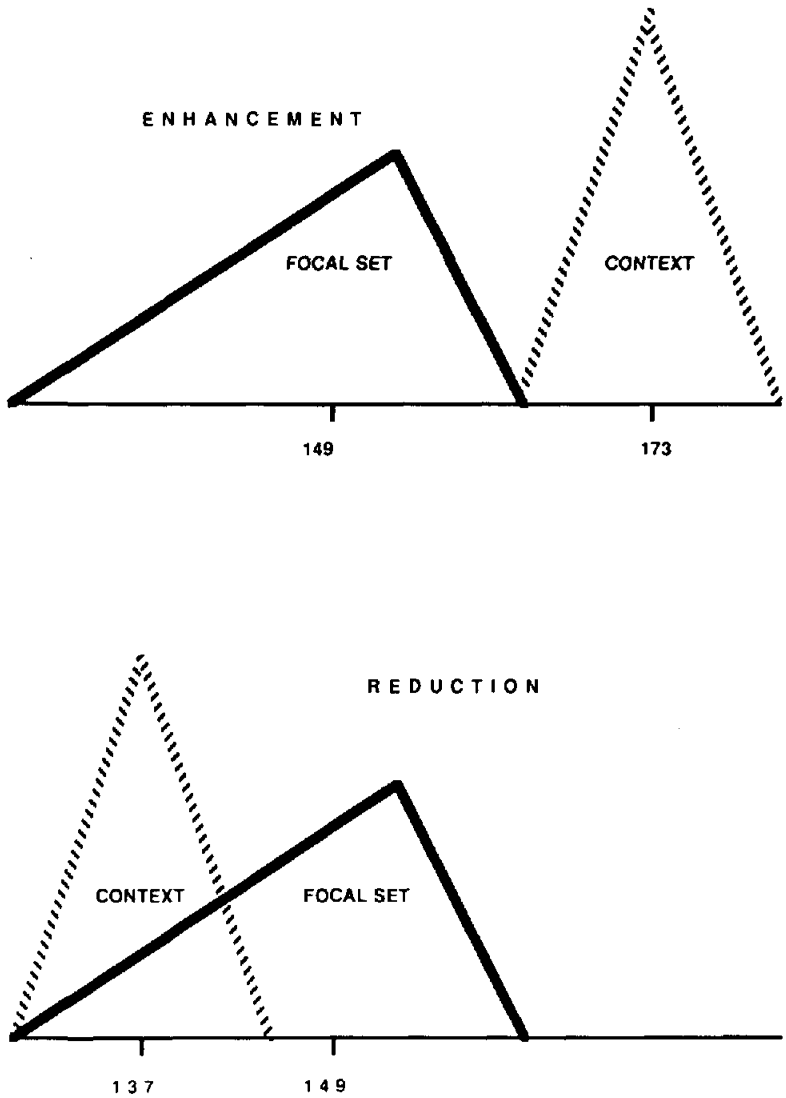

Figure 2. Schematic representation of the number distributions of the focal and contextual categories after Phase 2: Experiment 1.

condition (reduction vs. enhancement). As predicted, the mean estimates in the enhancement condition were more extreme than those in the reduction condition (lower when the direction of change was down and higher when the direction of change was up). The significant interaction term validates this pattern of results, $F(1, 52) = 5.54, p < .02$. Mean estimates were responsive to the presentation of the new, discrepant numbers presented in Phase 2. The significant main effect of direction indicates that estimates in Phase 2 were larger than estimates in Phase 1 when higher numbers were presented and smaller than estimates in Phase 1 when lower numbers were presented, $F(1, 52) = 283.49, p < .001$.

Although the changes in mean estimates in the focal category showed the predicted pattern, they did not reveal any information about perceived intercategory differences. If changes in the perceived means in the focal category were accompanied by equally large changes in the contextual category, the intercategory difference would have remained constant. Mathematically, the averages in the contextual categories remained constant throughout the entire experiment. To test whether subjects' perceptions remained stable as well, t tests for dependent measures were conducted for each group, comparing the perceived means in Phase 1 with the perceived means in Phase 2. In all six groups, the means were the same in both phases (all $ps > .07$). In the one case that showed a marginally significant change, the direction of the change was in the direction of greater perceived intercategory differentiation. When the mean of the focal category increased, the mean of the low contextual category decreased (Phase 1 M = 138.44; Phase 2 M = 136.88), $t(12) = 2.02, p < .07$.

Next we computed a difference measure of category change. The mean estimates in Phase 1 were subtracted from the mean estimates in Phase 2. Table 2 shows the estimated mean changes for the enhancement, reduction, and baseline conditions, as well as the averaged true change.

Table 1
Averaged Estimates of the Means of the Focal Category in Phase 1 and in Phase 2: Experiment 1

	Direction of change in focal category					
	Down			Up		
Condition	Reduction	Baseline	Enhancement	Reduction	Baseline	Enhancement
Phase 1						
Estimated *M*	154.00	155.39	155.75	155.59	154.92	154.36
SD	4.18	1.35	4.37	2.87	1.53	3.33
True *M*	155.50	155.50	155.50	155.50	155.50	155.50
Phase 2						
Estimated *M*	148.42	148.31	147.30	162.40	162.15	163.91
SD	3.36	2.91	4.60	2.46	3.65	4.80
True *M*	151.34	151.34	151.34	159.66	159.66	159.66

Note. The Direction–Down–Reduction and Direction–Up–Enhancement columns represent the conditions with a low contextual category. The Direction–Down–Enhancement and Direction–Up–Reduction columns represent the conditions with a high contextual category.

We analyzed the data of the two experimental conditions (enhancement and reduction) with a two-way analysis of variance (ANOVA), treating condition and direction as between-subjects variables with two levels each. As predicted by the category-accentuation hypothesis, greater absolute changes appeared in the enhancement condition than in the reduction condition. Specifically, we predicted an interaction between direction and condition because changes were either positive (direction up) or negative (direction down). This interaction was significant, $F(1, 53) = 8.21$, $p < .04$. Difference scores were negative when the category mean shifted downward and positive when the category mean shifted upward. The significant direction effect confirms this observation, $F(1, 53) = 246.11$, $p < .001$.

The averaged shifts in the two baseline conditions did not differ. Perceived category changes toward and away from the contextual category were of equal magnitude when the contextual category was far removed from the focal category. When averaged across direction, the shifts in the baseline conditions ($M = 7.15$) were only insignificantly greater than the shifts in the reduction conditions ($M = 6.20$), $t(46) = 1.06$, $p > .30$, and they differed with marginal significance from the shifts in the enhancement conditions ($M = 9.02$), $t(50) = 3.12$, $p < .09$. These results do not rule out the contribution of either inhibition of difference reduction or facilitation of difference enhancement to intercategory differentiation.

Table 2
Averaged Changes in Mean Estimates From Phase 1 to Phase 2: Experiment 1

	Direction of change in focal category			
	Down		Up	
Condition	*M*	*SD*	*M*	*SD*
Enhancement	−8.44	4.76	9.55	3.54
Reduction	−5.58	2.42	6.81	3.21
Baseline	−7.08	3.00	7.24	3.84
True change	−4.16		4.16	

It should be noted that, in every condition, the estimated changes were greater than the true change (all $ps < .001$). This overestimation may be attributed to a recency effect. Numbers that were shown in Phase 2 were more available in memory and were thus weighted more heavily than numbers shown in Phase 1.

Discussion

Overall, Experiment 1's findings cast a positive light on statistical intuition. Despite the challenging task of abstracting two cumulative means simultaneously, estimates were quite accurate, and in Phase 1 the estimated focal means were unaffected by the location of the contextual category. That is, there was no contrast effect. This finding conforms well with data from Malmi and Samson (1983). These authors presented two series of three-digit numbers whose means were either 1 point or 71 points apart. Each number was shown for either 1.5 s or 0.5 s. No contrast effects were observed unless the means were 1 point apart and the numbers were shown for only 0.5 s.

During Phase 2, estimates changed more when the focal mean moved away from the contextual mean than when it moved toward the contextual mean. In other words, perceived changes in intercategory differences were accentuated. This interpretation is supported by the fact that estimates in the contextual category remained constant across phases.

We suggested previously that category accentuation may arise from a two-stage process. Once subjects had learned to distinguish between the two categories, they subsequently bolstered this impression by giving greater weight to new elements that strengthened the perceived distinction than to elements that weakened it. However, as attractive as this explanation may be, there remains the possibility that a simple contrast effect can explain the data equally well.

We did not expect a contrast effect with unambiguous numerical stimuli, and in fact, at the end of Phase 1 the estimated means in the focal category did not vary as a function of the location of the context category. Specifically, the estimates were not displaced away from the contextual category. This finding alone, however, may not warrant the conclusion that contrast effects never occur. In Phase 1, the two distributions were adjacent and nonoverlapping. It may be, however, that contrast effects occur when two distributions overlap. If there is a general tendency to perceive categories as distinct, the presentation of overlapping distributions could in itself cause a contrast effect. Consider the stimuli presented in the reduction condition. At the end of Phase 2, the two distributions overlapped because the new information had increased the variance in the focal category. The perceived average of the focal distribution moved less toward the contextual category than it did in the enhancement condition. However, this finding may simply reflect a contrast effect, that is, a tendency to exaggerate the mean differences between two overlapping distributions. In Experiment 2, therefore, the same distributions as in Experiment 1 were presented, but now there was no distinction between the category-learning and category-change phases.

Experiment 2

Experiment 2 was a one-phase experiment in which the mean of the focal category was stable across all 12 blocks. In one condition, the focal and contextual categories overlapped throughout the experiment, rather than overlapping only during the second half of the experiment as in Experiment 1's reduction condition. In another condition, analogous to Experiment 1's enhancement condition, the two distributions never overlapped. According to the contrast hypothesis, the estimated mean in the focal category will vary with the location of the contextual category. That is, the mean estimates will be higher when the focal category is paired with a low contextual category than when it is paired with a high contextual category. According to the hypothesis of category accentuation, however, the estimated means in the focal category should not vary as a function of the location of the contextual category.

Method

Subjects. Forty-nine undergraduate students participated in this experiment in exchange for extra credit in their introductory psychology classes.

Procedure. The procedure in Experiment 2 was the same as in Experiment 1, with one critical difference: The sequence of presentation of numbers within the focal category was random, and thus there was no distinction between the two phases. For half the subjects, the mean of the focal category was 161.5 across estimation blocks (high focal). This category was analogous to the condition of upward change in Experiment 1. For the other half, the mean of the focal category was 149.5. This category was analogous to the condition of downward change in Experiment 1. Each of these focal categories was presented with a contextual category whose mean was either lower (M = 137.5) or higher (M = 173.5). The two categories overlapped when the high focal category was paired with the high contextual category or when the low focal category was paired with the low contextual category. The two distributions were distinct when the high focal category was paired with the low contextual category or when the low focal category was paired with the high contextual category. By the end of the experiment, the set of numbers presented in a given condition was identical to the set of numbers in the corresponding condition of Experiment 1. The resulting final distribution was described by the same parameters. As in Experiment 1, subjects estimated the overall mean of the focal and contextual categories 12 times.

Table 3
Averaged Estimates of Focal Category Means: Experiment 2

Focal category	Location of context category: High	Low	True M
High	161.19	160.82	161.50
SD	5.11	2.64	
Low	148.56	149.69	149.50
SD	4.09	4.54	
M	154.88	155.26	

Results

For each subject, the first 3 estimates were omitted from the analyses. We averaged and analyzed the following 9 estimates in a two-way ANOVA with context (high vs. low) and focal (high focal category vs. low focal category) as between-subjects variables. The four estimated means of the focal category and the two true means are shown in Table 3.

In all four cells, subjects' estimates were close to the true means. For the high focal category, subjects estimated significantly higher means than for the low focal category, $F(1, 45) = 213.77$, $p < .001$. The contrast hypothesis suggested that the estimated means would be larger in the presence of a low contextual category than in the presence of a high contextual category. The data did not support this prediction, $F(1, 45) < 1$. Similarly, estimates in the contextual category did not differ significantly as a function of the location of the focal category, $F(1, 45) < 1$.

Discussion

In Experiment 2, we showed that even when two distributions of numbers overlapped, the estimates of their respective means did not show contrast effects. Experiment 1's findings may therefore reflect the accentuation of category change that depends on the learning of intercategory differences and the subsequent differential weighting of new information that reduces or enhances this learned distinction.

What is the bearing of these findings on the learning and modification of social categories? Experiment 1 demonstrated category accentuation with numerical stimuli. Judgments about people who form social groups are typically based on stimuli that are more complex. Social perceivers, forming impressions about people or behaviors, may hold a priori expectations about the two groups. Do such expectations or stereotypes affect the way in which group means are abstracted from relevant person information? We address this question in Experiment 3.

Experiment 3

We repeated the experimental procedure of Experiment 1, with the numerical stimuli representing a measure of an attri-

Table 4
Averaged Estimates of the Means of the Focal Category in Phase 1 and in Phase 2: Experiment 3

	Direction of change in focal category			
	Down		Up	
Condition	Reduction	Enhancement	Reduction	Enhancement
Phase 1				
Estimated *M*	153.26	155.70	154.87	155.25
SD	5.11	2.43	2.27	2.58
True *M*	155.50	155.50	155.50	155.50
Phase 2				
Estimated *M*	148.65	146.21	162.93	164.04
SD	4.37	2.86	2.90	3.29
True *M*	151.34	151.34	159.66	159.66

Note. The Direction-Down–Reduction and Direction-Up–Enhancement columns represent the conditions with a low contextual category. The Direction-Down–Enhancement and Direction-Up–Reduction columns represent the conditions with a high contextual category.

bute for members of two groups of people. The mean difference between the two categories reflected a social stereotype. The aim of this experiment was to test whether the category-accentuation effect appears with social categories.

Method

Subjects. Forty-two undergraduate students who were enrolled in introductory psychology classes participated for extra course credit.

Procedure. As stated, the procedures were identical to those of Experiment 1; however, subjects were told that the numbers represented the body weight (in pounds) of marathoners and sprinters. Subjects were also explicitly told that sprinters are on the average heavier than marathoners. Thus, during the category-learning phase, subjects received information confirming the induced expectation: Lower weights were always associated with marathoners and higher weights were always associated with sprinters. Again, during the category-change phase, new information in the focal group rendered the overall mean more similar or more dissimilar to the mean of the contextual group. That is, the subjects' expectations were confirmed when the intercategory difference was enhanced, but were disconfirmed when the intercategory difference was reduced. Before a number appeared on the screen, a prompt was shown for 500 m, reading either *Marathoner* or *Sprinter.* We omitted the baseline condition because it was not possible to construct a plausible contextual category with a mean weight sufficiently different from that in the focal category.

Results

As in the previous experiments, estimates of the mean weight of the marathoners and the sprinters showed no contrast effect at the end of the category-learning phase. The estimated mean of the focal category was not greater in the presence of a low contextual category ($M = 154.26$) than in the presence of a high contextual category ($M = 155.29$). Table 4 shows the mean estimates and standard deviations for the two phases in the four experimental conditions.

As a test of the category-accentuation hypothesis, we performed a two-way ANCOVA on the average estimates in Phase 2, with direction and condition as between-subjects variables and the average estimates in Phase 1 as a covariate. The predicted two-way interaction was only marginally significant, $F(1, 37) = 3.23$, $p < .08$. Nevertheless, the pattern of means, as summarized in Table 4, is highly consistent with the results in Experiment 1, showing a trend for estimates in the enhancement condition to be more displaced from the estimates in Phase 1 than were estimates in the reduction condition. Compared with Phase 1, mean estimates in Phase 2 changed in the direction predicted by the new numbers presented, $F(1, 37) = 238.79$, $p < .001$. Once again, mean estimates in the contextual category did not vary from Phase 1 to Phase 2 (all $ps > .10$).

We computed the difference measure of category change by subtracting estimates in Phase 1 from estimates in Phase 2. The results for the two experimental groups are shown in Table 5.

The category-accentuation hypothesis predicted an interaction between the direction of change and the enhancement and reduction conditions. This interaction was significant, $F(1, 38) = 4.34$, $p < .04$, in a two-way ANOVA. The absolute size of the difference scores was greater in the enhancement condition than in the reduction condition. There was also the expected effect for direction, indicating that the difference scores became either positive or negative, depending on whether smaller or greater numbers were added to the focal category, $F(1, 38) = 138.61$, $p < .001$. No other effect was significant. As in the first experiment, perceived change in the enhancement condition ($M = 9.14$) was significantly greater than the true mean change ($M = 4.36$), $t(20) = 6.21$, $p < .001$. Perceived change in the reduction condition ($M = 6.34$) exceeded the true change only marginally, $t(19) = 2.02$, $p < .06$.

Discussion

Experiment 3 shows that the accentuation effect in category change appears with judgments of physical attributes of two groups of people. However, the induced expectation of weight differences between marathoners and sprinters did not increase category accentuation beyond the effect observed in Experiment 1. It is possible that subjects disregarded prior beliefs when forming impressions about groups. Beliefs about the association between athletic category and body weight may have been only tenuous in comparison with the readily available information about individual group members. Studies on the effect of stereotypes on judgments of individual people have shown that individuating information can override beliefs

Table 5
Averaged Changes in Mean Estimates From Phase 1 to Phase 2: Experiment 3

	Direction of change in focal category			
	Down		Up	
Condition	*M*	*SD*	*M*	*SD*
Enhancement	−9.48	4.38	8.79	3.41
Reduction	−4.61	5.13	8.07	4.45
True change	−4.16		4.16	

about the category, particularly when the stereotypic beliefs are weak (Krueger & Rothbart, 1988; Locksley, Borgida, Brekke, & Hepburn, 1980). It is possible that subjects did not hold sufficiently strong beliefs about the relative weight of runners to produce an accentuation effect greater than the effect with numbers alone.

General Discussion

We draw three conclusions from these experiments. First, judgments of the central tendency of category attributes change more when intercategory differences are enhanced than when they are reduced. Second, category accentuation is obtained with numbers alone as well as with numbers that represent person attributes. Third, category accentuation is independent of perceptual contrast effects. Because the use of unambiguous stimuli precluded perceptual distortion, higher-order mental processes need to be invoked to explain category accentuation.

Explaining Category Accentuation

What are the specific processes that cause category accentuation? Although the mean estimation task required thinking, not feeling, motivational factors cannot be ruled out. Gordon Allport suspected that a "principle of least effort inclines us to hold coarse and early generalizations as long as they can possibly be made to serve our purposes" (Allport, 1954, p. 172). If people adopt the principle of least effort, they should prefer easily discriminable categories, because this separation should facilitate inferences from category to exemplar and vice versa. The motivational hypothesis could be tested by offering subjects rewards to motivate them for accuracy. Tajfel and Wilkes (1963) speculated that accentuation effects may be eliminated by monetary incentives. However, there is no experimental research pertaining to this hypothesis. Informal questioning after our experiments revealed that, contrary to the motivational explanation, most subjects were concerned about the accuracy of their judgments. Although they had been explicitly instructed not to think of the experiment as a test of mathematical ability, many subjects were more interested in how they did than in the theoretical debriefing.

An alternative to the motivational hypothesis is the view that a lack of cognitive capacity interfered with statistical intuitions. Dawes (1964) found illogical simplifications in memory for set relations. When two partially overlapping or "disjunctive" sets were presented ("Some A are B and some A are not B"), subjects were liable to remember Set A as completely overlapping with or included in Set B ("All A are B"). This error of "overgeneralization" was more frequent than the opposite error of "pseudodiscrimination." The relation "All A are B" (inclusion) was rarely misremembered as "Some A are B and some A are not B." Note that overgeneralization reduces complexity, whereas pseudodiscrimination increases it. Category accentuation, like overgeneralization, reduces complexity by turning a disjunctive set relation into a simpler, exclusive relation. Either technique reflects Dawes's view of "people as one-bit minds in a two-bit world" (Dawes, 1964, p. 457).

Which memory processes might be involved in the simplification of perceived category structure? According to the well-documented von Restorff effect, information is better remembered if it is distinct from other information that is presented at the same time (von Restorff, 1932). The category-modifying numbers in the enhancement condition, in comparison with those in the reduction condition, were extreme and unique stimuli. They were extreme because they differed more from the grand mean of all presented numbers, and they were unique because they appeared only as elements of one category.

Despite its appealing cogency and simplicity, the validity of the memory hypothesis is questionable. If selective retrieval of extreme numbers were the main cause for biased average estimation, then we should have found evidence for category accentuation in Experiment 2. In Experiment 2, half the subjects were presented with overlapping distributions throughout the experiment, not only during Phase 2. The finding that category accentuation occurred only when category change was preceded by a phase of category learning (Experiments 1 and 3) may be the result of a combination of motivational and cognitive factors. An extensive cognitive and social literature has gathered evidence for confirmatory biases that lend stability to beliefs (cf. Nisbett & Ross, 1980). Confirming information, for example, is encoded more effortlessly and recalled better than disconfirming information (Rothbart, Evans, & Fulero, 1979). The studies reported here can be interpreted as cases of belief verification. After subjects had learned that two categories of numbers were distinct, new information that would reduce the perceived intercategory difference was disconfirming, whereas information that would maintain or enhance the perceived difference was confirming. The strong version of the argument of belief maintenance does not hold. If belief perpetuation were the supreme purpose of cognition, any category changing information should meet with resistance and neglect. The exact magnitude of the initially learned intercategory difference would have to be protected. Neither a greater nor a smaller difference would be accepted. Still, a more complex version of the same argument seems tenable. Category changing information is resisted only if the very existence of a difference is jeopardized. The magnification of intercategory differences, on the other hand, is comparatively inconsequential.

Modifiability of Social Beliefs

The primary motivation for this research was to explore processes involved in the stability of social beliefs and, in particular, social stereotypes. In Experiments 1 and 3, perceived category change was greater than true change; this seems to contradict Allport's claim that categories are "stubborn and resistant to change" (Allport, 1954, p. 17). However, the present experiments are only a special case of category change because all the category-modifying elements for any given subject were presented late in the learning procedure, and all the elements modified the mean in the same direction. Comparisons between enhancements and reductions of intercategory differences were made between subjects. Contact with members of social groups, on the other hand, may yield information that both enhances and reduces perceived group differences. The same observer may witness the increase of the group variance at both ends of the distribution. To resurrect a classic example from Katz and Braly (1933), consider a person who holds a stereotype that Germans are stolid. Further assume that the ascription of this trait implies comparisons with other relevant ethnic groups or between Germans and humanity in general (cf. McCauley & Stitt, 1978). In a contact situation with members of the stereotyped group, the observer is likely to encounter a sample of individuals who vary on the trait. Some individuals may appear exceedingly stolid; others may appear emotional. Thus, there is information that both enhances and reduces perceived differences between the target group and the comparison group. Although in an actuarial sense all the new information may average out to the initial perception of the average stolidness of Germans, the category-accentuation effect suggests that perceived intercategory differences become greater and produce a stronger stereotype. The implications of this research for stereotype change are still speculative. Future research will show whether the mere increase in variance will alter the perceived central tendency of a category.

References

Allen, V. L., & Wilder, D. A. (1979). Group categorization and attribution of belief similarity. *Small Group Behavior, 10,* 73–80.

Allport, G. W. (1954). *On the nature of prejudice.* Reading, MA: Addison-Wesley.

Beach, L. R., & Swensson, R. G. (1966). Intuitive estimation of means. *Psychonomic Science, 5,* 161–162.

Billig, M., & Tajfel, H. (1973). Social categorization and similarity in intergroup behavior. *European Journal of Social Psychology, 3,* 27–52.

Busemeyer, J. R., & Myung, I. J. (1988). A new method for investigating prototype learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 3–11.

Dawes, R. M. (1964). Cognitive distortion. *Psychological Reports, 14,* 443–459.

Dawes, R. M., Singer, D., & Lemons, F. (1972). An experimental analysis of the contrast effect and its implications for intergroup communication and the indirect assessment of attitude. *Journal of Personality and Social Psychology, 21,* 281–295.

Eiser, J. R. (1971). Enhancement of contrast in the absolute judgment of attitude statements. *Journal of Personality and Social Psychology, 17,* 1–10.

Hovland, C. I., Harvey, O. J., & Sherif, M. (1957). Assimilation and contrast effects in reactions to communication and attitude change. *Journal of Abnormal and Social Psychology, 55,* 244–252.

Katz, D. (1960). The functional approach to the study of attitudes. *Public Opinion Quarterly, 24,* 163–204.

Katz, D., & Braly, K. W. (1933). Racial stereotypes of one hundred college students. *Journal of Abnormal and Social Psychology, 28,* 280–290.

Krueger, J., & Rothbart, M. (1988). Use of category and individuating information in the prediction of personality attributes. *Journal of Personality and Social Psychology, 55,* 187–195.

Locksley, A., Borgida, E., Brekke, N., & Hepburn, C. (1980). Sex stereotypes and social judgment. *Journal of Personality and Social Psychology, 39,* 821–831.

Malmi, R. A., & Samson, D. J. (1983). Intuitive averaging of categorized numerical stimuli. *Journal of Verbal Learning and Verbal Behavior, 22,* 547–559.

McCauley, C., & Stitt, C. L. (1978). An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology, 36,* 929–940.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings.* Englewood Cliffs, NJ: Prentice-Hall.

Rabbie, J. M., & Horwitz, M. (1969). Arousal of ingroup–outgroup bias by chance win or loss. *Journal of Personality and Social Psychology, 13,* 269–277.

Rothbart, M., Evans, M., & Fulero, S. (1979). Recall for confirming events: Memory processes and the maintenance of social stereotypes. *Journal of Experimental Social Psychology, 15,* 343–355.

Spencer, J. A. (1963). A further study of estimating averages. *Ergonomics, 6,* 255–256.

Tajfel, H., & Wilkes, A. L. (1963). Classification and quantitative judgment. *British Journal of Social Psychology, 54,* 101–114.

Von Restorff, H. (1932). Ueber die Wirkung von Bereichsbildungen im Spurenfeld [On the effects of category formation in the field of memory traces]. In W. Koehler & H. von Restorff (Eds.), *Analyse von Vorgaengen im Spurenfeld: 1. Psychologische Forschung, 18,* 299–342.

Wilder, D. A., & Thompson, J. E. (1988). Assimilation and contrast effects in the judgments of groups. *Journal of Personality and Social Psychology, 54,* 62–73.

Received March 18, 1988
Revision received December 12, 1988
Accepted December 19, 1988 ■