UC Davis

UC Davis Previously Published Works

Title

Development and mental representation of stereotypes.

Permalink

<https://escholarship.org/uc/item/93w872pk>

Journal

Journal of personality and social psychology, 70(6)

ISSN

0022-3514

Author

Sherman, JW

Publication Date

1996-06-01

DOI

10.1037//0022-3514.70.6.1126

Peer reviewed

Development and Mental Representation of Stereotypes

Jeffrey W. Sherman
Northwestern University

A mixed model of stereotype representation was tested. Experiment 1 examined the development of stereotypes about novel groups. Results showed that, at low levels of experience, stereotypic group knowledge is derived from information about particular group exemplars. However, as experience increases, an abstract group stereotype is formed that is stored and retrieved independently of the exemplars on which it was based. Results of Experiment 2 suggest that preexisting stereotypes about well-known groups are represented as abstract structures in memory. These results indicate that stereotypical knowledge is most likely to be exemplar-based in the absence of abstract stereotypes. The implications of these findings for other aspects of stereotyping and social perception are discussed.

The stereotype has been a central construct in the field of social psychology since Lippman (1922) first coined the term in his influential analysis of intergroup perception. For many years research on stereotypes focused almost exclusively on identifying the content of stereotypes about different social groups (Brigham, 1971; Gilbert, 1951; Karlins, Coffman, & Walters, 1969; Katz & Braly, 1933). More recently, researchers have begun to examine some of the specific processes by which stereotypes exert their influence on social perception, judgment, and behavior (for a review see Hamilton & Sherman, 1994). Confirming the importance of the construct, this research has shown that the activation of a stereotype can affect all aspects of social information processing, including attentional allocation (Bodenhausen, 1988), behavioral interpretation (Darley & Gross, 1983; Kunda & Sherman-Williams, 1992; Sagar & Schofield, 1980), inference making (Bodenhausen & Wyer, 1985; Krueger & Rothbart, 1988), and retrieval (Bodenhausen & Lichtenstein, 1987; Cohen, 1981; Hamilton & Rose, 1980; Snyder & Uranowitz, 1978). An activated stereotype can also influence the types of information perceivers seek about targets in the first place (Kunda, 1990; Skov & Sherman, 1986; Snyder & Swann, 1978) and can direct the behavior of perceivers in ways that lead to stereotype confirmation, pro-

ducing self-fulfilling hypotheses (Snyder, Tanke, & Berscheid, 1977; Word, Zanna, & Cooper, 1974).

But what does it mean to activate a stereotype? What exactly gets activated? More precisely, what is a stereotype? The latter question can be answered on at least two different levels. On one level, a stereotype can be defined as “a cognitive structure that contains the perceiver’s knowledge, beliefs, and expectations about a human group” (Hamilton & Trolie, 1986, p. 133). According to this view, when a stereotype is activated one’s knowledge, beliefs, and expectations about a group are brought to mind. This general conceptualization is widely accepted among social psychologists and provides an adequate level of description for many purposes.

However, this definition is representationally vague. It does not specify the precise nature of the “cognitive structures” that contain stereotypical information. How are one’s knowledge, beliefs, and expectations about a group represented in memory? On a more specific level, then, a stereotype may be defined in terms of the particular kinds of mental representations that form the basis of one’s knowledge about social groups.

Defining stereotypes at the representational level is crucial because different kinds of representations have different implications for other important questions about stereotypes. How are stereotypes retrieved and applied in social perception? How do stereotypes develop over time? What function do stereotypes serve in social cognition? The answers to these questions (and many others pertaining to stereotype change, stereotype flexibility, subtyping, stereotype assessment, and the role of categorization in social perception) depend in large part on how a stereotype is represented in memory (Smith, 1990, 1992). However, despite the attention social psychologists have paid to stereotypes, very little research has attempted to define stereotypes at the representational level.

Representational Models of Stereotyping

Recent research on the mental representation of both social (e.g., Allen & Ebbesen, 1981; Carlston, 1980; Carlston & Skowronski, 1986; Klein, Loftus, Trafton, & Fuhrman, 1992; Park, 1986; Sherman & Klein, 1994; Smith & Zarate, 1990) and non-

This research was supported by an American Psychological Association Dissertation Research Award; a Graduate Division Dissertation Fellowship from the University of California, Santa Barbara; and a General Affiliates Graduate Dissertation Fellowship from the University of California, Santa Barbara.

Parts of this research were included in my doctoral dissertation submitted to the University of California, Santa Barbara. I would like to thank my committee: Dave Hamilton, Stan Klein, Diane Mackie, and Faith Gleicher, for greatly improving the quality of this research. I also thank Jim Sherman, Steve Stroessner, Eliot Smith, and Chick Judd for their thoughtful comments on drafts of this article and Leigh Frost and Elisabeth Schmidt for their assistance in conducting this research.

Correspondence concerning this article should be addressed to Jeffrey W. Sherman, Department of Psychology, Northwestern University, 2029 Sheridan Road, Evanston, Illinois 60208-2710. Electronic mail may be sent via the Internet to sherm@nwu.edu.

social (e.g., Busemeyer, Dewey, & Medin, 1984; Elio & Anderson, 1981; Homa, Dunbar, & Nohre, 1991; Homa, Sterling, & Trepel, 1981; Malt, 1989) concepts has been influenced by the distinction between abstract and exemplar-based knowledge. For a given concept, abstract knowledge consists of a summary representation of the typical features of the concept that either has been abstracted from experience with multiple exemplars of the concept or has been learned from outside sources (e.g., Posner & Keele, 1968; Rosch, 1975). In contrast, exemplar knowledge consists of separate representations of the concept's specific known exemplars in memory (e.g., Brooks, 1978; Hintzman, 1986). This distinction forms the basis for the following discussion of representational models of stereotypes.

Pure Abstraction Models of Stereotypes

Traditionally, social psychologists have adhered to pure abstraction models of stereotype representation. Lippman (1922) defined stereotypes as oversimplified generalizations about categories of people (see also Brigham, 1971). More recently, these generalizations have been labeled *schemata* (e.g., Taylor & Crocker, 1981), *prototypes* (e.g., Brewer, Dull, & Lui, 1981; Cantor & Mischel, 1979), *expectancies* (e.g., Hamilton, Sherman, & Ruvalo, 1990), and *Bayesian base rates* (e.g., McCauley & Stitt, 1978), among others. Although there are minor differences between these models, they all conceptualize stereotypes as abstract summaries of the typical features of a social group. These abstractions develop as perceivers acquire information about the group (e.g., Posner & Keele, 1968). This information may be acquired from firsthand personal experience with group members or through social learning from family, friends, and the media. Although stereotypes may be derived from episodes involving particular group exemplars, once formed, the stereotypes are stored separately from those exemplars in memory. According to this view, the stereotype is an autonomous representation that is independently stored and retrieved for future purposes. Once a target has been categorized as a member of a particular social group, that group's stereotype may be activated and applied toward the perception of that target. Although perceivers may be able to retrieve particular exemplars, the use of these exemplars does not constitute stereotyping.

Pure Exemplar Models of Stereotypes

Recently, Smith (1990, 1992; Smith & Zarate, 1992) challenged the notion of abstraction-based stereotypes, arguing that group exemplars must play an important role in stereotyping (see also Linville, Fischer, & Salovey, 1989). Smith based his model on instance-based models of categorization (Hintzman, 1986; Nosofsky, 1987). The primary basis for this challenge is the notion that abstraction-based stereotypes are simply too inflexible to account for phenomena such as perceived group variability, extensive subtyping, and context effects in stereotyping (for reviews see Hamilton & Sherman, 1994; Smith, 1990). According to Smith's model, stereotypes do not exist as independently stored knowledge structures. Rather, ad hoc "stereotypes" are created by activating particular instances of group members and summarizing their features. Social targets insti-

gate this process by acting as retrieval cues that recruit exemplars stored in memory. Once a group summary has been formed, it may be stored for future use. However, upon activation it will be weighted as merely another instance of the group to be combined with other instances in a novel summarization to be applied to the target.

Whereas categorization processes play a central role in pure abstraction models of stereotypes (social categorization necessarily precedes stereotype activation), they do not in pure exemplar models. According to pure exemplar models, social perception is dependent on the set of exemplars that are activated by the target. For any given target, a large variety of exemplars are brought to bear on the impression formation process. These activated exemplars do not all necessarily belong to the same social categories as each other or the target. Rather, the exemplars activated are simply those that share the most features in common with the target (depending on attentional factors; Smith & Zarate, 1992). Upon activation, the attributes of these exemplars are summarized to form expectancies, inferences, and judgments about the target. Stereotyping occurs to the extent that the activated exemplars all belong to a particular social category. For example, if a perceiver is forming an impression of a Black professor, both *Black* exemplars and *professor* exemplars (as well as many others) will be activated. The proportion of these different types of exemplars activated depends on which attributes of the target are attended to most carefully. If the *Blackness* attribute of the target receives particular attention and is salient, then a large proportion of the activated exemplars will be Black exemplars, and the person will be perceived as stereotypically Black. If *professorly* attributes are salient, a large proportion of professor exemplars will be activated, and the person will be perceived as a stereotypical professor.

Pure exemplar models permit a good deal of flexibility in stereotyping. Any time stereotyping occurs, a novel "stereotype" must be re-created by activating particular group exemplars. Therefore, the content of a "stereotype" is constantly changing, depending on the target and the set of exemplars activated. Indeed, because any given "stereotype" is created dually by the particular target involved and the particular set of activated exemplars, "stereotypes" will not generalize to novel targets. Even if novel targets are very similar to the target that created an initial group "stereotype," the new targets will stimulate their own exemplar activation and summation processes. Although it may be very similar to the previously created summarization, a new "stereotype" will be created for every new stimulus. Thus, the Black "stereotype" applied to our Black professor example will be (at least) slightly different each time it is created, even if the process occurs twice in 15 min. According to this view, there is no such thing as a stored abstract group representation that is independently retrieved and used. Thus, the model proposes that although stereotyping occurs through activated exemplars, stereotypes do not exist as independent mental representations.

Although pure exemplar models of stereotyping may help to account for the specificity and contextuality of social cognition, they are not without their limitations. Pure exemplar models have been challenged as inefficient (e.g., Hamilton & Mackie, 1990; Hamilton & Sherman, 1994; Nosofsky, Palmeri, & McKimley, 1994) and unable to explain the occurrence of on-line (e.g., Hastie & Park, 1986; Park & Hastie, 1987) and abstrac-

tion-based (e.g., Lingle & Ostrom, 1979; Sherman & Klein, 1994) judgments. Others have argued that exemplars must be joined together by some sort of category definition or theory of inclusion criterion (Medin & Wattenmaker, 1987). Otherwise, when asked to make a category judgment, perceivers could not activate appropriate exemplars.

Previous Research

Very little research has directly investigated the mental representation of stereotypes. Those experiments that have examined this issue have relied largely on correlational data. For example, a number of studies have demonstrated a lack of a positive correlation between the information that participants can remember about group members and the judgments they make about the group as a whole (Hamilton, Dugan, & Trolie, 1985; Judd & Park, 1988; McConnell, Sherman, & Hamilton, 1994a, 1994b; Park & Hastie, 1987; Stroessner, Hamilton, & Mackie, 1992). This suggests that group knowledge is not based on information about particular group members, supporting an abstraction-based model of stereotype representation.

Evidence provided by Park and Hastie (1987) also supports an abstraction-based model of stereotypes. Their data demonstrated that group behaviors that had been made salient through repetition were remembered better by participants but did not exert any additional influence on participants' judgments about the group. This suggests that participants' judgments were not based on what they could remember about the group but were instead based on abstract group impressions. However, it is possible that participants recognized that some of the behaviors had been repeated and therefore accordingly discounted those behaviors' impact on group judgments.

Other research lends support to exemplar-based models of stereotype representation. In contrast to the aforementioned research, a number of experiments have demonstrated a positive correlation between the information recalled about a group and judgments made about the group, suggesting exemplar-based group knowledge (Fiedler, Russer, & Gramm, 1993; Hamilton et al., 1985; Mackie, Sherman, & Worth, 1993; Manis & Paskewitz, 1987; Pryor, 1986; Rothbart, Fulero, Jensen, Howard, & Birrell, 1978). However, correlations between recall and judgment are difficult to interpret. For instance, it is possible that these correlations occur because participants use their judgments as retrieval cues for group members. Furthermore, it could be that both the judgments and the recall are driven by participants' reliance on a stored group stereotype. If this were the case, the stereotype could bias recall by focusing participants on particular group members during encoding as well as recall (Hamilton & Sherman, 1994).

Further support for exemplar models has been provided by computer simulations reported by Smith (1991) and Linville et al. (1989). Smith's simulation demonstrated that an exemplar-retrieval model can account for the formation of illusory correlations between social groups and the attributes that describe them (e.g., Hamilton & Sherman, 1989). Linville et al.'s simulation showed that an exemplar-based model was sufficient to account for participants' judgments that in-groups possess more intragroup variability than out-groups (the out-group homogeneity effect). Although such simulations may demonstrate

the *sufficiency* of exemplar activation in producing various judgmental biases, they do not demonstrate that such exemplar-based processes *necessarily* occur when perceivers make their judgments.

Other recent evidence also seems to support an exemplar-based model of stereotypes. Schwarz and Bless (1992) showed that participants' judgments about politicians in general were influenced by the particular politicians that participants had previously been asked to think about. For instance, participants who were asked to think about politicians involved in a political scandal subsequently rated politicians in general to be less trustworthy than did participants who did not think of the scandalous politicians. Similarly, Bodenhausen, Schwarz, Bless, and Wanke (1995) demonstrated that participants' perceptions of racism toward Blacks were influenced by the nature of the Black exemplars (e.g., Michael Jordan and Jesse Jackson) that had previously been made salient. These experiments demonstrated that, once activated, exemplars may indeed influence group judgments. However, these results did not demonstrate that such exemplars are spontaneously activated by perceivers when they think about these social groups. Moreover, the results are moot about whether perceivers also possess and apply abstract stereotypes in addition to the salient exemplars in the judgment. Thus, these experiments demonstrated that activated group exemplars may affect group judgments, but they do not shed light on the underlying mental representation of group knowledge.

A study often cited as providing support for exemplar-based processing was conducted by Lewicki (1985). Participants expected a newly encountered person with short hair to be unfriendly simply because of a previous encounter with an unfriendly person who also happened to have short hair. It appears that participants accessed the first person as a basis for their judgment of the second person. Smith (1990) argued that an abstraction-based model of stereotyping could not account for the impact of such single instances on judgments. However, it is unlikely that participants ever had stereotypes about the friendliness of people with short hair.

Finally, Smith and Zarate (1990) demonstrated that categorization judgments about novel social categories tend to be dominated by exemplar-based processes. After learning about four or five members of two social groups (Group A and Group B), participants were asked to classify new individuals as being from Group A or Group B. Results demonstrated that classification judgments tended to be based on the target's similarity to specific exemplars from the two groups and not on the target's similarity to the groups' prototypes. Interestingly, the only situation that produced prototype-based judgments was when participants were given the group prototype before they learned about the group members. Thus, the presence of a preexisting stereotype instigated abstraction-based classifications.

A Mixed Model of Stereotypes

In light of the theoretical and empirical limitations of both pure abstraction and pure exemplar models, many researchers have adopted mixed models of representation that contain both abstract and exemplar information (Busemeyer et al., 1984; Carlston, 1980; Carlston & Skowronski, 1986; Elio & Ander-

son, 1981; Homa et al., 1981, 1991; Klein et al., 1992; Malt, 1989; Nosofsky et al., 1994; Park, 1986; Ross, Perkins, & Tenpenny, 1990; Sherman & Klein, 1994; Smith & Zarate, 1990). According to these models, both abstract and exemplar representations may form the basis of social knowledge under different conditions. Given the difficulty that both pure abstraction and pure exemplar models have in accounting for the accumulated data on stereotype representation, such an approach would seem useful in addressing this issue.

One important variable that affects perceivers' reliance on exemplars and abstractions is the amount of experience perceivers have with the target to be judged. During the initial stages of learning about a target, judgments are based on the activation of particular exemplars because too few exemplars have been encountered to form useful abstract knowledge. However, as the number of exemplars encountered increases, an abstract representation of the target evolves, which then serves as the basis for subsequent judgments (e.g., Homa et al., 1991; Klein & Loftus, 1993b; Klein et al., 1992; Ross et al., 1990; Sherman & Klein, 1994). Thus, exemplar use decreases as target experience increases.

Research on impression formation (Klein et al., 1992; Park, 1986; Sherman & Klein, 1994) and self-perception (Klein & Loftus, 1993a; Klein et al., 1992) has supported such a mixed model of social perception. These studies indicate that knowledge of the attributes that describe other individuals and the self may be represented as either exemplar or abstract knowledge, depending on the perceiver's degree of experience with the other individual and the degree of self-experience within particular contexts. At low levels of experience, knowledge is based on activated exemplar behaviors pertaining to the other person and the self. At high levels of experience, abstract impressions are formed, and judgments need not rely on exemplar activation.

Although there is now considerable support for such mixed models of nonsocial knowledge and knowledge about other individuals and the self, the role of experience has never been examined in research on stereotype representation. The goal of the present research was to test such a mixed model of stereotypical knowledge. According to this model, the mental representation of stereotypic attributes depends on the perceiver's degree of experience with the group in question. Initially, when few group exemplars have been encountered, knowledge of the typical features of the group will be computed on the basis of the activation of particular group exemplars. If a "stereotype" of the group is to be applied, it must be computed from activated exemplars. However, with sufficient experience with group members (or secondhand accounts of their attributes), perceivers will form abstract representations of the attributes that are stereotypical of the group. Once formed, these abstract group impressions (stereotypes) become the basis for subsequent stereotyping processes. Thus, stereotypic knowledge is exemplar-based primarily when no abstract stereotype exists. However, when relevant abstract stereotypes exist, they will form the basis for stereotyping independently of stored exemplars. This model acknowledges the contribution of exemplar-based processes in intergroup perception without discarding the abstract stereotype construct altogether.

The results reported by Lewicki (1985) and Smith and Zarate (1990) are consistent with the mixed model of stereotyping

that has been proposed. Both studies demonstrated a tendency toward exemplar-based processes when relevant stereotypes had not been strongly established. The goal of the present research was to more directly examine the mental representation of stereotypic knowledge and how that knowledge develops as experience with a group increases.

Experiment 1

To determine the extent to which group knowledge is based on the retrieval of group exemplars, the present research used a priming paradigm called the *task facilitation paradigm* (e.g., Klein, Loftus, & Burton, 1989). The task facilitation paradigm is based on the following logic. Suppose participants perform two tasks in succession. If participants utilize information in performing the first task that is necessary for performing the second task, then the time required to perform the second task should be less than if that information had not been made available (Collins & Quillian, 1970; Klein et al., 1992; Macht & O'Brien, 1980). Thus, examining the degree to which performing one task facilitates the performance of a second task is a way to determine the extent to which two tasks rely on the same information. As the overlap in the information used in performing the two tasks increases, the time necessary to perform the second task should decrease.

The specific task facilitation manipulation used in these experiments involved three different tasks. The *describes* task required participants to decide whether a stimulus trait described in general a group of people presented to the participants (e.g., "Does the word *kind* describe the group?"). The *recall* task required participants to retrieve from memory a specific behavioral incident in which a member of the target group behaved in accordance with the stimulus trait (e.g., "Remember a specific incident in which a member of the group behaved in a *kind* manner?"). The *define* task required participants to generate a definition for the stimulus trait (e.g., "Think of the meaning of the word *kind*"). A trial consisted of performing two of these tasks in succession, an initial task and a target task, on the same trait word.

In the present experiments, the target task in a trial was always a recall task. Sometimes the initial task was a describes task, and sometimes it was a define task. The primary dependent measure was the amount of time participants needed to perform the recall target task, given that they had previously performed an initial describes or define task. If knowledge about the typical features of a group is based on information about group exemplars, then it should take less time to recall a group behavior following a describes task than following a define task. This is because activating the group impression to perform the describes task means activating group exemplars whereas generating a definition does not. Therefore, a describes task should facilitate a subsequent recall task because retrieving an exemplar should take less time if exemplars have recently been activated. Past research has supported the assumption that the define task acts as a control condition in which behavioral exemplars are not activated (Klein & Loftus, 1993b; Klein et al., 1992).

In contrast, if knowledge about the typical features of a group consists of abstractions and not exemplars, then performing an

initial describes group judgment task should not facilitate exemplar retrieval relative to an initial define task. This is because activating the group impression to perform the describes task would not involve the activation of behavioral exemplars.

Overview

Experiment 1 examined the development and the mental representation of stereotypic knowledge about a group for which participants had no preexisting stereotype. After learning either a relatively small or a relatively large amount of information about members of an unspecified group of college students, participants performed either a describes-recall or a define-recall task facilitation trial.

Predictions

Pure abstraction model. The pure abstraction model predicts that group exemplars should at no time form the basis for group judgments. Therefore, the time required to perform the recall target task should not differ as a function of the initial task. This should be true whether participants have received a small or a large amount of information about the group. At no time does a group description task involve the spontaneous activation of exemplars.

Pure exemplar model. The pure exemplar model predicts that group judgments are always based on activated group exemplars. Therefore, the time required to perform the exemplar retrieval task should be shorter following an initial describes task than following an initial define task. Because performing the describes task involves the spontaneous activation of group exemplars, the subsequent recall task should take relatively little time compared with when it follows a define task. This difference should hold regardless of whether participants have received a small or a large amount of information about the group. Because abstract knowledge does not develop over time according to pure exemplar models, exemplar summarization forms the necessary basis for group judgments regardless of the perceiver's level of experience with the group.

Mixed model. The mixed model proposes that exemplars will form the basis of group knowledge only when abstract group knowledge has yet to be created. Therefore, the mixed model predicts that participants who receive a small amount of information about the group will base their judgments of the group on activated exemplars. However, participants who receive a large amount of information about the group should form abstract group impressions (stereotypes) of the group's features. These abstractions, and not specific exemplars, would then form the basis for participants' judgments about the group. Therefore, for participants in the low-experience condition, recall-task response times should be faster following an initial describes task than following an initial define task because exemplars are activated in order to make the group judgment. However, for participants in the high-experience condition, recall-task response times should be equal following the describes and define tasks because the group judgment need not rely on activated exemplars; abstract knowledge has been created.

Method

Participants. The 156 participants were recruited from the Northwestern University participant pool and received partial course credit for their participation. Participants were tested in groups of 1–6.

Materials and design. Participants read either one block or four blocks of information about a target group of people. Each block of information contained two different kind behaviors (e.g., "Stopped to let another car into the line of traffic"), two different intelligent behaviors (e.g., "Studies photography in his spare time"), two behaviors that did not imply either kindness or intelligence (e.g., "Took a walk around the block after dinner"), and four demographic items (e.g., "Was born in Phoenix, Arizona"). On scales from 0 (*not at all*) to 10 (*very*), the intelligent behaviors were rated as moderately intelligent ($M = 6.83$), and the kind behaviors were rated as moderately kind ($M = 7.42$).¹ The assignment of stimulus items to the two experience conditions (one block and four blocks) was randomly determined. Two different one-block and four-block stimulus sequences were created. Participants performed the initial task (describes or define) and the recall target task in reference to either *kind* or *intelligent*. In summary, the experiment was a 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (trait: kind vs. intelligent) $\times$ 2 (stimulus replication: Order A vs. Order B) between-subjects design.

Procedure. Participants were told that they would be reading a series of descriptions that had been provided by a group of students "who all belong to the same club at a large Midwestern university." These descriptions were ostensibly drawn from a larger pool of information provided by the group. Participants were told that each item presented had been provided by a different member of the group. However, no names were presented in association with any of the stimulus items. Participants were asked to form an impression of "what the group is like in general" as they read about the group. A microcomputer presented the descriptions, one every 6 s.

After reading the stimuli, participants were trained to perform the describes, define, and recall tasks. During training, participants were asked to think of a group of their friends as they performed the describes and recall tasks. Each practice trial consisted of performing two tasks in succession, an initial task and a target task, on a trait unrelated to kindness or intelligence. For the describes task, participants decided whether the trait described their group of friends; for the recall task, participants recalled a behavioral incident in which one of their friends manifested the trait; and for the define task, participants generated a definition for the trait. Participants performed six different combinations of initial task (describes, recall, or define) and target task (describes, recall, or define).

Upon completion of the practice trials, participants performed a single test trial using the group of students they had read about as the target group. The trial began with one of two cues for the initial task appearing on the screen: "DESCRIBES STUDENT GROUP" (describes task) or "DEFINE" (define task). A stimulus trait (kind or intelligent) ap-

¹ To decrease the likelihood that participants would spontaneously draw trait inferences about the group on the basis of the first behavior, the chosen stimulus behaviors were only moderately indicative of the traits they represented (e.g., Winter & Uleman, 1984). Theoretically, the less trait-prototypical the behaviors are, the more evidence should be required to form abstract impressions (e.g., Buss & Craik, 1983, 1984; Trope, 1986; Trope, Cohen, & Alfieri, 1991). Indeed, Sherman and Klein (1994) found that abstract impressions of individuals developed almost immediately with the presentation of highly, but not moderately, diagnostic behaviors. Although highly diagnostic behaviors are frequently used in studies of impression formation (e.g., Dreben, Fiske, & Hastie, 1979; Srull, 1981), more moderate behaviors would seem to be more common in everyday life.

peared beneath the task cue 2 s later. The cue and the stimulus trait remained on the screen until the participant indicated that he or she had completed the initial task by pressing the space bar.² At this point, the initial task was removed from the screen. After a 2-s pause, the cue for the recall target task ("RECALL STUDENT GROUP") appeared on the screen above the same stimulus trait, and a timer started in the computer. Again, the cue and the trait remained on the screen until the participant pressed the space bar to indicate completion of the recall task. With this response, the timer stopped, the participant's recall task latency was recorded, and the trial ended.

Immediately following the test trial, participants were asked to write down the specific kind or intelligent behavior they recalled when performing the recall target task. The data of 14 participants who failed to report a stimulus behavior were removed from the data set.³

Results

The data of participants with response latencies greater than 2.5 standard deviations above the mean were excluded from the data set. This resulted in the removal of data for 4 additional participants. Therefore, the analyses were based on the data from 138 participants. For purposes of data normalization, all analyses were based on a log transformation of the response latencies. All means are reported in milliseconds.

A 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (trait: kind vs. intelligent) between-subjects analysis of covariance (ANCOVA) was conducted on the recall target task response latencies.⁴ An average of each participant's response latencies on the practice recall tasks was used as the covariate. The analysis produced a significant main effect for amount of information, such that recall latencies were shorter in the four-block condition ($M = 4,920$) than in the one-block condition ($M = 5,973$), $F(1, 129) = 5.44$, $p < .05$. However, this effect was qualified by the predicted two-way interaction between amount of information and type of initial task, $F(1, 129) = 4.01$, $p < .05$. As predicted, participants in the one-block condition took less time to recall a behavior following an initial describes task ($M = 5,360$) than following an initial define task ($M = 6,586$), $t(64) = 1.93$, $p < .05$, one-tailed. In contrast, participants in the four-block condition recalled behaviors equally quickly following an initial describes task ($M = 4,985$) and an initial define task ($M = 4,855$), $t < 1$. These recall-task response latencies are presented in Figure 1.⁵

Discussion

The recall-task response latencies provided strong initial support for a mixed model of stereotype representation. In the one-block condition, an initial group judgment task facilitated the subsequent retrieval of a group exemplar behavior as compared with an initial define task. This demonstrates that participants accessed specific group exemplars from memory in order to judge the group. However, in the four-block condition, participants took equally long to retrieve a group exemplar following a describes and a define task. These participants did not access exemplars in order to make judgments about the group.

These results indicate that the mental representation of knowledge about the typical features of a group changes as experience with a group increases. At low levels of experience, group typicality knowledge is represented in terms of specific

group members. However, as experience increases, an abstract group impression (a stereotype) is formed that is stored and retrieved independently of knowledge about particular group members. Contrary to the predictions of the pure abstraction model, group typicality knowledge was exemplar-based in the one-block condition. Contrary to the predictions of the pure exemplar model, group typicality knowledge was not exemplar-based in the four-block condition.

Experiment 2

Experiment 2 examined more closely the mixed model's predictions regarding the mental representation of stereotypes.

² Participants were not asked to report their responses during the experimental trials; rather, they were instructed to generate responses to the tasks in their heads. Past research has corroborated this approach. First, studies by Klein and Loftus (1993b) and Klein et al. (1992) demonstrated that response latencies for performing the describes task on highly and lowly self- and mother-descriptive traits were relatively fast compared with the time required to perform the task on moderately descriptive traits. This inverted-U function replicated the results of several studies in which participants reported their self- and other-descriptive judgments at the time they were made (e.g., Kuiper, 1981; Lord, Gilbert, & Stanley, 1982). As expected, participants performed the mental describes task more easily for traits that were most clearly related (positively and negatively) to the target's personality. Furthermore, latencies to recall, but not to report, specific trait-relevant behaviors were monotonically related to the level of descriptiveness of the trait involved. As trait-descriptiveness decreased, recall latencies increased (Klein & Loftus, 1993b; Klein et al., 1992). These data replicated the findings of Klein and Loftus (1991), who asked participants to report recalled behaviors. These data suggest that participants mentally performed the experimental tasks in the manner requested.

³ This nonreport rate of 9% is very close to that reported by Sherman and Klein (1994) in their experiment on individual impression formation (nonreport rate of 11%). A 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (trait: kind vs. intelligent) between-subjects analysis of variance was conducted to determine whether reporting failures were distributed evenly across experimental conditions. This analysis revealed no significant effects, demonstrating that such failures were distributed evenly across the conditions.

⁴ The stimulus replication variable had no effect and interacted with no variables in Experiment 1 or Experiment 2. Therefore, all analyses were collapsed across this variable in both experiments.

⁵ It is worth noting that the crucial comparisons are those within a particular level of behavioral experience. In the one-block condition, it was predicted that describes judgments would be based on activated behaviors. Therefore, recall should be faster following an initial describes task than following an initial define task. However, in the four-block condition, it was predicted that describes judgments would not be based on exemplars. As a result, it should take equally long to recall a behavior following the two initial tasks. However, the fact that recall should be facilitated following a describes task (relative to the define control task) in the one-block condition, but not in the four-block condition, does not imply that recall following a describes task should be faster in the one-block condition than in the four-block condition. Any advantage in recall times afforded by exemplar activation in the one-block condition may be offset by a general tendency for recall to require less time as behavioral experience increases, as exhibited by the amount of information main effect (see also Klein & Loftus, 1991; Klein et al., 1992; Myers, O'Brien, Balota, & Toyofuku, 1984; Park, 1989; Sherman & Klein, 1994).

Figure 1. Mean recall target task latency as a function of initial task and level of familiarity in Experiment 1.

The data from Experiment 1 demonstrated that, although group knowledge may be exemplar-based when experience with a group is minimal, perceivers do develop, store, and use abstract stereotypes as experience increases. However, one potential criticism of Experiment 1 is that the experimentally induced abstract impressions that participants developed in the four-block condition do not correspond to real-world stereotypes, which tend to be quite stable and are usually based on a good deal of experience, firsthand or otherwise. Experiment 2 addressed this issue by examining the mental representation of two well-known real-world stereotypes.

Another goal of Experiment 2 was to examine an alternative explanation for the data from Experiment 1. Keenan (1993) argued that findings such as those obtained in the four-block condition of Experiment 1 would not necessarily rule out exemplar-based models of stereotyping. Keenan suggested that judgments in both the one-block and four-block conditions might be exemplar-based. Recall facilitation occurs in the one-block, but not the four-block, condition because of a fan effect. According to this argument, when participants judge the group in the one-block condition, two trait-relevant behaviors are activated, and the degree of activation that spreads to each of the two items is sufficient to produce facilitation on the subsequent recall task. However, for judgments in the four-block condition, activation must spread to eight trait-relevant behaviors. Thus, each item receives relatively little activation, resulting in no demonstrable recall facilitation (Anderson & Bower, 1973). Thus, Keenan argued that, even though judgments in both the one-block and four-block conditions are exemplar-based, only the one-block condition produces facilitation. This possibility was further examined in Experiment 2.

Overview

The experimental stimuli, methods, procedure, and independent and dependent variables were identical to those used in Experiment 1, with one exception. Whereas the target group in Experiment 1 was an ill-defined group for which participants possessed no preexisting stereotype, the target group in Experiment 2 was a more cohesive social group for which people hold a common stereotype. Half the participants were told that they would be reading about a group of engineers, and half were told that they would be reading about a group of priests. These groups were chosen because they are stereotypically intelligent and kind, respectively. After reading one block or four blocks of information about the group, participants performed the same task facilitation trials as in Experiment 1, using the specific group of engineers or priests they had read about as the target. The initial task (describes or define) and the target recall task in a trial pertained to either a stereotype-consistent trait (engineer-intelligent or priest-kind) or a stereotype-irrelevant trait (engineer-kind or priest-intelligent). By examining participants' recall task latencies, it was possible to chart the influence of participants' stereotypes of engineers and priests on their judgments about the group and, thus, directly examine the mental representation of these stereotypes.

Predictions

Pure abstraction model. The pure abstraction model predicts that all knowledge about the typical features of a group is represented abstractly. Therefore, target recall task latencies should not be facilitated by an initial describes judgment task. This should be true regardless of the amount of information

participants have learned about the group and regardless of whether the trait judged is consistent with or irrelevant to the group stereotype.

Pure exemplar model. The pure exemplar model predicts that group typicality knowledge is always based on group exemplars. Therefore, target recall task latencies should be facilitated by an initial describes task in both the one-block and four-block conditions. This should be true regardless of whether the trait judged is stereotype-consistent or not. The knowledge that the group consists of engineers or priests should in no way alter the manner in which judgments are made about the group.

Mixed model. The mixed model proposes that the stereotype associated with the category label will provide participants with abstract trait knowledge about the target group. This trait knowledge will form the basis for an abstract stereotypical impression of the target group, which will be used to make judgments about the group pertaining to stereotype-consistent traits. As a result, stereotype-consistent group judgments should not be based on activated exemplars, even in the one-block condition.

When no preexisting stereotype of the target group exists, judgments made in the absence of induced abstract group impressions (the one-block condition of Experiment 1) are based on the activation of group exemplars. However, providing participants with a preexisting group stereotype should attenuate the use of exemplar-based strategies by providing the participants with ready-made abstract group impressions. Therefore, even if the perceiver has yet to induce abstract knowledge about the group from the presented information, exemplar activation will be unnecessary for group judgments. Rather, the preexisting stereotype can form the basis for an abstract stereotypical impression, precluding the necessity of exemplar use.

Importantly, this should be true only for stereotype-consistent traits. Thus, for the group of engineers, intelligence judgments should not rely on exemplar activation in the one-block or the four-block condition. Recall-task response times should be equal following the define and describe tasks in both the one-block and four-block conditions. The category label (engineers) provides access to abstract information about the group's intelligence, allowing perceivers to immediately form abstract stereotypical impressions. Therefore, group intelligence judgments do not require exemplar activation, even in the one-block condition.

However, judgments about the kindness of the engineers should follow the same pattern as that in Experiment 1. Because the stereotype of engineers includes no information about their kindness, knowledge about the target group's kindness should follow the developmental pattern found in Experiment 1. Initially, in the one-block condition, kind judgments will be based on activated exemplar information. However, by the end of the fourth block, participants should develop abstract impressions of the target group's kindness, and judgments will no longer rely on exemplar activation. Of course, the opposite pattern of trait results would be predicted for the group of priests. Thus, the results for the stereotype-irrelevant trait provide a conceptual replication for the pattern of stereotype development observed in Experiment 1.

Fan effect hypothesis. If Keenan (1993) is correct, the preexisting stereotypes presented to participants in Experiment 2

should produce results identical to those in Experiment 1. Once again, the one-block, but not the four-block, condition should produce recall facilitation. This should be true regardless of whether the trait to be judged is stereotype-consistent or stereotype-irrelevant. This contrasts with the mixed model's prediction of no facilitation in either the one-block or the four-block condition for stereotype-consistent judgments. The mixed model suggests that the preexisting stereotype will form the basis of an abstract impression that will be used in lieu of exemplars in the one-block condition.

Method

Participants. The 156 participants were recruited from various undergraduate classes at the University of California, Santa Barbara. Some participants ($n = 105$) were paid \$5 for their participation, and others received partial course credit ($n = 51$). Participants were tested in groups of 1–6.

Materials and design. The stimulus materials were identical to those used in Experiment 1. Participants read either one block or four blocks of information about a target group of people. Participants were told that the group they were reading about was either a group of engineers or a group of priests. Two different stimulus presentation orders were created, with items being randomly assigned to the one-block and four-block conditions. Participants performed the initial task (describes or define) and the recall target task in reference to either a stereotype-consistent trait (engineer–intelligent or priest–kind) or a stereotype-irrelevant trait (engineer–kind or priest–intelligent).⁶ In summary, the experiment was a 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (target group: engineer vs. priest) $\times$ 2 (trait: kind vs. intelligent) $\times$ 2 (stimulus replication: Order A vs. Order B) $\times$ 2 (type of participant: paid vs. unpaid) between-subjects design.

Procedure. Participants were told that they would be reading a series of descriptions that had been provided by a group of engineers or priests. These descriptions were ostensibly obtained through a large survey conducted on a wide variety of different occupational groups. Participants were told that the items they would be reading about the group had been drawn from a larger pool of information provided by the group. They were also led to believe that each item presented had been provided by a different member of the group. No names were associated with any of the information. Participants were asked to form an impression of "what this group of engineers (priests) is like in general" as they read about the group. A microcomputer presented the descriptions in one of two random orders, one every 6 s.

After reading the stimuli, participants were trained to perform the describes, define, and recall tasks, as described in Experiment 1. Upon completion of the practice trials, participants performed a single test

⁶ In a free-response pretest, 7 out of 11 participants offered the trait *kind* as typical of priests, but none offered the trait *intelligent*. In a separate free-response pretest, 77 out of 99 participants offered the trait *intelligent* as typical of engineers, but none offered the trait *kind*. An additional pretest was conducted in which participants were asked to rate either "priests in general" or "engineers in general" on a number of different traits, including kindness and intelligence, on scales ranging from 1 (*not at all*) to 10 (*very*). An analysis of these ratings produced a significant two-way interaction, such that priests were rated as more kind ($M = 8.08$) than engineers ($M = 6.81$) and engineers were rated as more intelligent ($M = 9.43$) than priests ($M = 7.18$), $F(1, 38) = 28.41$, $p < .01$. These data also indicated that priests were perceived to be more kind than intelligent, whereas engineers were perceived to be more intelligent than kind.

trial using the specific group of engineers or priests they had read about as the target group. The trial was identical to that used in Experiment 1, except for slight changes in the task cues that appeared on the computer screen. The cues appearing on the screen in Experiment 2 were "DESCRIBES GROUP" (describes task), "DEFINE" (define task), and "RECALL GROUP" (recall task). The instructions stressed that participants' describes and recall task responses should pertain to "the group of engineers/priests you read about in the first part of this experiment."

Once again, immediately following the test trial, participants were asked to write down the specific kind or intelligent behavior they recalled when performing the recall target task. The data of 14 participants who failed to report a stimulus behavior were removed from the data set.⁷

Results

The data of participants with response latencies greater than 2.5 standard deviations above the mean were excluded from the data set. This resulted in the removal of data for 3 additional participants. Therefore, the analyses were based on the data from 139 participants. For purposes of data normalization, all analyses were based on a log transformation of the response latencies. All means are reported in milliseconds.

For clarification purposes (and given the theoretical importance of stereotype consistency in the predictions), the target type (priest vs. engineer) and trait type (kind vs. intelligent) variables were collapsed into one variable in the data analyses. This variable divided participants into two groups: one for those who performed their tasks in reference to stereotype-consistent traits (priest-kind and engineer-intelligent), and one for those who performed their tasks in reference to stereotype-irrelevant traits (priest-intelligent and engineer-kind). A 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (consistency of trait rated: consistent vs. irrelevant) $\times$ 2 (type of participant: paid vs. unpaid) between-subjects ANCOVA was conducted on the recall target task response latencies. Once again, an average of each participant's response latencies on the practice recall tasks was used as the covariate.

This analysis produced a significant main effect for initial task type, with participants recalling a behavior more quickly following an initial describes task ($M = 5,120$) than following an initial define task ($M = 6,843$), $F(1, 119) = 6.68$, $p < .05$. This effect was qualified by a two-way interaction between initial task type and trait consistency, which showed that the tendency for recall latencies to be faster following a describes task than following a define task was true only for stereotype-irrelevant, and not stereotype-consistent, traits, $F(1, 119) = 3.83$, $p = .0527$. Participants recalled a stereotype-irrelevant behavior more quickly following an initial describes task ($M = 4,627$) than following a define task ($M = 7,979$), $t(67) = 2.51$, $p < .05$. In contrast, stereotype-consistent behaviors were recalled equally quickly following both kinds of tasks (describes $M = 5,614$ and define $M = 5,706$), $t < 1$. This interaction is depicted in Figure 2.⁸

These results indicate that group exemplars were activated for judgments about stereotype-irrelevant but not stereotype-consistent traits. However, the predicted three-way interaction between initial task type, trait consistency, and amount of in-

formation was not significant. Unexpectedly, judgments about stereotype-irrelevant traits were exemplar-based in both the one-block and four-block conditions. Possible explanations for this result are discussed later in the *Development of stereotype-irrelevant impressions* section.

The ANCOVA also produced a significant two-way interaction between trait consistency and amount of information, which showed that, collapsed across initial task type, it took longer to recall stereotype-irrelevant behaviors in the one-block condition ($M = 7,755$) than in the four-block condition ($M = 4,851$) but took an equal amount of time to recall stereotype-consistent behaviors in the one-block ($M = 5,497$) and four-block ($M = 5,823$) conditions, $F(1, 119) = 4.43$, $p < .05$.

Discussion

Stereotype representation. The results of Experiment 2 support the mixed model's predictions regarding stereotype representation. Whereas stereotype-irrelevant recall latencies were facilitated by an initial describes judgment task, stereotype-consistent latencies were not. This finding demonstrates that participants accessed specific group exemplars in order to judge the group only when their group stereotype was irrelevant to the judgment task. When participants were able to apply a relevant stereotype to their judgment task, they did not need to activate exemplars in order to make the judgment. These data are consistent with the mixed model's assertion that the group labels and their associated stereotypes formed the basis for abstract group impressions that were used in making stereotype-consistent judgments. This suggests that stereotypes about engineers and priests are represented as abstract knowledge.

The fact that stereotype-consistent judgments in the one-block condition were not exemplar-based is particularly telling. The results from Experiment 1 and from the stereotype-irrelevant condition of Experiment 2 suggest that participants were unable to form abstract group impressions on the basis of one block of information. As a result, judgments in the one-block condition were based on activated exemplars. Only when participants possessed an applicable stereotype were they able to make group judgments in the one-block condition without activating exemplars. Because participants were unable to induce abstract knowledge from the presented information in the one-

⁷ This yielded a failure rate of 9%. A 2 (amount of information: one block vs. four blocks) $\times$ 2 (initial task: describes vs. define) $\times$ 2 (target group: engineer vs. priest) $\times$ 2 (trait rated: kind vs. intelligent) $\times$ 2 (stimulus replication: Order A vs. Order B) between-subjects analysis of variance was conducted to determine whether reporting failures were distributed evenly across experimental conditions. This analysis yielded no significant results, demonstrating that such failures were evenly distributed across conditions.

⁸ An additional 2 (initial task: describes vs. define) $\times$ 2 (consistency of trait rated: consistent vs. irrelevant) $\times$ 2 (target type: engineer vs. priest) between-subjects ANCOVA was conducted to ensure that the initial task type by trait consistency interaction generalized across both target groups. This analysis produced a marginally significant initial task type by trait consistency interaction, $F(1, 127) = 2.74$, $p = .10$, which was not qualified by the target factor ($F < 1$ for the three-way interaction), indicating that the interaction generalized across both targets.

Figure 2. Mean recall target task latency as a function of initial task and trait-stereotype consistency in Experiment 2.

block condition, the trait knowledge used by participants in the stereotype-consistent condition would have had to come from elsewhere. It would seem that the stereotype provided by the group label supplied participants with abstract trait knowledge and, thus, abstract impressions that could be relied on for group judgments.

These results bolster the conclusions drawn from Experiment 1. In Experiment 1, it was suggested that the abstract impressions participants developed in the four-block condition may not have corresponded to real-world stereotypes. However, in Experiment 2, existing stereotypes were studied directly.

Fan effect hypothesis. The results from Experiment 2 also indicate that a fan effect cannot explain the data from Experiment 1. If a fan effect were responsible for the results of Experiment 1, then Experiment 2 should have produced the same results: recall facilitation in the one-block condition but not the four-block condition. However, in Experiment 2 stereotype-consistent judgments yielded no facilitation in either the one-block or the four-block condition. In contrast, stereotype-irrelevant judgments produced facilitation in both the one-block and four-block conditions. Clearly, both of these results are incompatible with a fan effect account of the data.

Development of stereotype-irrelevant impressions. It was hypothesized that recall latencies for participants in the stereotype-irrelevant conditions would replicate the results from Experiment 1. At low levels of experience, stereotype-irrelevant judgments were predicted to be exemplar-based. However, by the fourth block of information, participants were expected to have formed abstract impressions of the group, which would then form the basis for group judgments. As predicted, stereotype-irrelevant latencies in the one-block condition were facilitated by an initial describes task, indicating that judgments in this condition were exemplar-based. This replicates the data

from Experiment 1, indicating that, at low levels of experience, judgments about a group for which a stereotype cannot be applied are based on activated exemplars.

Unexpectedly, stereotype-irrelevant recall latencies were also facilitated by an initial describes task in the four-block condition, indicating that judgments in this condition were also exemplar-based. It appears that participants in this condition were unable to induce abstract group knowledge, even after receiving four blocks of information. In Experiment 1, no stereotypes were provided about the stimulus group, and participants were able to form abstract group knowledge by the fourth block of information. This suggests that the stereotype provided in Experiment 2 may have actually inhibited the development of abstract knowledge about the stereotype-irrelevant trait.

Similar findings have been reported by Chapman and Chapman (1969) in their classic studies on the illusory correlation. In one study, participants perceived positive correlations between patient types (homosexuals) and stereotypical responses to Rorschach tests (e.g., those having to do with feminine clothing), although no such correlations were present in the stimuli. This is the classic illusory correlation. In another study, participants continued to perceive illusory correlations between patient types and stereotypical responses even when a true correlation existed in the stimuli between patient type and nonstereotypical responses. In this study, participants did not perceive the true correlation between patient type and nonstereotypical responses. However, in another study in which the stimuli contained no stereotypical responses, participants were able to accurately identify true correlations between patient type and nonstereotypical responses. These results suggest that the presence of stereotypical responses prevented participants from perceiving true correlations between patient type and nonstereotypical responses. As in the present research, the stereotype

seemed to inhibit the development of stereotype-irrelevant knowledge.

It is intriguing to consider that this may be yet another means by which stereotypes perpetuate themselves. In addition to the various consistency biases on attention, encoding, retrieval, and behavior, stereotypes may preserve themselves by interfering with perceivers' ability to form impressions of the group that are unrelated to the domain of the stereotype. This would help to maintain the stereotype by ensuring that stereotypical attributes continued to occupy a central position in impressions of the group. Relative to other attributes, stereotypic knowledge would be (among other things) more abstract, more accessible, more extreme, and more easily activated and applied to social perception. This is a potentially important effect of stereotypes that warrants further investigation.

Which exemplars were retrieved? There is a potential alternative explanation for the results in Experiment 2. It could be argued that group judgments in both the stereotype-consistent and stereotype-irrelevant conditions were exemplar-based but that the exemplars on which the stereotype-consistent group judgments were based were not those presented in the stimulus materials. Rather, participants may have derived their stereotypic responses from memories of other engineers and priests stored from prior real-world experiences that were then applied to the judgment task. However, this seems implausible for a variety of reasons. First, it seems very unlikely that the exemplars most directly relevant to the judgment at hand would not be accessed during this process. If the target probe ("the group of engineers/priests you read about") were to activate any exemplars, it would be those exemplars most similar to the probe (members of the group) that would most likely be activated. Even if outside exemplars were accessed, particular group members should have been accessed as well, according to pure exemplar models.

In addition, analyses of two other dependent measures collected in Experiment 2 argue against this possibility. First, participants' initial describes task response latencies were measured. An analysis of these latencies indicated that trait judgments were made more quickly in the four-block condition ($M = 3,464$) than in the one-block condition ($M = 4,467$) $t(66) = 2.90, p < .05$. In contrast, define task responses (the control condition) were made equally quickly in the four-block ($M = 3,516$) and one-block ($M = 3,672$) conditions, $t < 1$. This produced a significant initial task type by amount of information interaction, $F(1, 121) = 4.07, p < .05$. This interaction was present in both the stereotype-consistent and stereotype-irrelevant conditions. Such decreases in stereotype-consistent judgment times are inconsistent with an extraexperimental exemplar account of the data. The latencies of judgments based on the activation of outside exemplars should not be affected by the number of experimental exemplars encountered.

Furthermore, at the end of Experiment 2, participants were asked to rate how kind and intelligent the stimulus group was on a scale from 1 (*not at all*) to 6 (*very*). An analysis of these ratings indicated that participants made more extreme judgments about the groups in the four-block condition ($M = 4.99$) than in the one-block condition ($M = 4.69$), $F(1, 137) = 8.21, p < .01$. This effect was present for both the consistent and irrelevant traits. These results also argue against an extraexperi-

mental exemplar account of the primary data. If nonstimulus exemplars formed the basis for the participants' group judgments, then trait ratings should not have differed as a function of amount of information presented. The number of stimulus group members presented (one block or four blocks) would be irrelevant to a group judgment based on outside exemplars.

The mixed model, however, can account for both of these findings. In contrast to pure exemplar models, the mixed model permits the on-line development of impressions. These on-line encoding processes are not disengaged once abstract stereotypical impressions are formed (Sherman & Klein, 1994). As additional stereotypical behaviors are encountered, the impression may be accessed and updated with the new behaviors. As a result, the stereotypical impression may become more accessible (leading to faster judgment latencies) and extreme as consistent information accumulates. Finally, an extraexperimental exemplar account cannot account for the data from Experiment 1, which demonstrated the development of abstract group stereotypes.

General Discussion

This research attempted to define stereotypes at the representational level. The results of two experiments supported a mixed representational model of stereotypical knowledge. Experiment 1 demonstrated that the mental representation of stereotypical knowledge depends on the perceiver's degree of experience with the group. At low levels of experience, knowledge of the typical features of the group was based on particular group exemplars stored in memory. However, as experience increased, an abstract stereotype was induced that became the basis for subsequent judgments about the group.

Experiment 2 further supported the conclusion that abstract stereotypes are stored independently in memory and may be retrieved as the basis for group judgments. In Experiment 2, the relevance of a well-developed stereotype to participants' judgments was directly manipulated. The results demonstrated that knowledge about the typical features of the group was exemplar-based only when participants did not possess an applicable group stereotype. Thus, the stereotype label provided access to trait knowledge about the group's stereotype-consistent traits independently of group exemplars. The results suggest that this trait knowledge formed the basis for abstract impressions that were then used to make stereotypic judgments about the group. These findings run counter to the predictions of both pure abstraction and pure exemplar models of stereotyping.

Implications for Social Perception

Social categorization. The results of these experiments have important implications for many issues in social perception. First, they confirm the importance of the categorization process in stereotyping. The results of Experiment 1 and the irrelevant condition of Experiment 2 indicate that participants were not able to induce abstract impressions of the group on the basis of one block of information. Then how were participants in the one-block condition of Experiment 2 able to make stereotype-consistent judgments without reference to exemplars? The answer is that the occupational category labels presented to

participants provided direct access to knowledge about stereotype-consistent traits. It would seem that this knowledge formed the basis for abstract group impressions that were used to make stereotypical judgments, rendering exemplar summarization unnecessary. Importantly, the presentation of group labels in itself did not diminish exemplar-based processing. Only those judgments that were consistent with the category stereotype benefited in this way. This confirms the central role of social categorization in the application of stereotypical knowledge. Contrary to the predictions of pure exemplar models, these results show that specific stereotypical knowledge is directly associated with category labels (e.g., engineer-intelligent and priest-kind). Upon categorization, a target may be understood in terms of its category stereotype.

Subtyping and stereotype flexibility. These results also have implications for subtyping processes and stereotype flexibility. Smith (1990) argued that abstract stereotypes are too inflexible to account for the number of group subtypes that people possess and the contextual flexibility inherent in group judgments. However, the data from these experiments highlight the important distinction between group *judgments* and group *stereotypes*. It is certainly correct to point out that, in many conditions, group judgments may be exemplar-based. For instance, judgments about previously unconsidered subtypes (women with brown eyes who work as certified public accountants) are most likely exemplar-based. However, such cases do not imply that people do not possess abstract stereotypes about groups with which they have more experience (e.g., priests and engineers). Furthermore, this is not to say that people do not develop abstract stereotypes about frequently encountered subgroups (e.g., grandmotherly types; Brewer et al., 1981).

Similarly, different contexts may elicit the activation of different exemplars when the intelligence of an engineer one has just met is considered (e.g., the engineer looks exactly like Albert Einstein or Dan Quayle). These activated exemplars may very well affect the judgment arrived at by the perceiver (Bodenhausen et al., 1995; Schwarz & Bless, 1992). However, this does not imply a lack of abstract stereotypes about engineers. Those exemplars are most likely activated in addition to, not instead of, an abstract group stereotype. To the extent that an abstract stereotype has been activated, it will affect the relative intelligence assigned to the engineer. The judgment may be higher or lower depending on who the engineer looks like, but it will always be higher than a rating of the intelligence of a fast-food cook who also looks like Einstein or Quayle.

Thus, an important conclusion from this research is that it is possible to acknowledge the impact of exemplars on social perception without having to claim that abstract stereotypes do not exist. One important topic for future research is to identify those situations where exemplars are activated in addition to, or even instead of, an abstract stereotype. This issue will be further addressed below.

Stereotype change. Smith (1990) also argued that stereotypes change too readily to be abstraction-based. Whereas exemplar-based knowledge should change frequently as new exemplars are encountered, abstraction-based knowledge should change slowly, as disconfirming information is averaged into the abstraction (Smith, 1990). The previously discussed work of

Lewicki (1985) is often cited as evidence that stereotypes are easily changed. However, as has already been mentioned, it seems rather unlikely that participants in Lewicki's experiment ever possessed stereotypes about the friendliness of people with short hair. Therefore, this analysis again fails to note the distinction between judgments and stereotypes. It is perfectly reasonable to acknowledge the impact of exemplars in Lewicki's experiment and still claim that abstract stereotypes exist. Certainly, knowledge about a nonstereotyped group may change rapidly (as in Experiment 1 of the present research).

Stereotype function. One important function of abstract stereotypes is to increase cognitive efficiency. Through the act of categorization, abstract stereotypes reduce the amount of information to which perceivers must attend. Social stimuli that have been grouped together can be treated as functionally equivalent, reducing the need to form individualized impressions of each category member (Allport, 1954; Brewer, 1988; Fiske & Neuberg, 1990; Hamilton & Sherman, 1994; Lippman, 1922). Abstract stereotypes also increase cognitive efficiency by expanding the base of knowledge that perceivers may apply toward social perception. By providing useful expectancies about group members' personalities and behavior, abstract stereotypes allow the perceiver to go beyond the information given (Allport, 1954; Hamilton & Sherman, 1994; Medin, 1988; Taylor, 1981). On the basis of a target's group membership and accompanying abstract stereotype, a perceiver may infer the target's personality attributes without having to carefully attend to the target's behavior (Brewer, 1988; Fiske & Neuberg, 1990), preserving resources that may be applied to other mental tasks (e.g., Macrae, Milne, & Bodenhausen, 1994).

The results of Experiments 1 and 2 suggest that abstract stereotypes not only reduce perceivers' reliance on individuating information, but also reduce their reliance on exemplar-based processing. In Experiment 1, group judgments in the one-block condition were exemplar-based. However, by the fourth block of information, an abstract stereotype of the group had been formed, and exemplars were no longer accessed for group judgments. The results of Experiment 2 are even more telling. When participants possessed both group exemplars and a stereotypical group label (providing access to trait knowledge about the consistent traits), the exemplars were not used for group judgments. This was true even in the one-block condition, which had been shown to elicit exemplar-based judgments when stereotypic knowledge was not available (Experiment 1 and the irrelevant condition of Experiment 2). These data suggest that basing an impression on an abstract stereotype is more efficient than forming both individuated impressions based on target behavior and inferences based on analogy to stored exemplars (e.g., Homa et al., 1981, 1991; Klein et al., 1992; Sherman & Klein, 1994). In addition to being more efficient, perceivers may also find abstract knowledge to be more diagnostic and reliable than exemplar knowledge.

Theoretical Issues

Source of group knowledge. The recall target task response latencies demonstrated that exemplars did not form the basis for group judgments in the four-block condition of Experiment 1 or in the stereotype-consistent conditions of Experiment 2.

The mixed model argues that the judgments were instead based on abstract group impressions that were either induced from the stimulus exemplars (Experiment 1) or provided directly by the relevant stereotypes (Experiment 2). However, such conclusions cannot be made with absolute certainty. The task facilitation method was not designed to assess the use or nonuse of abstract trait impressions. The use of such knowledge may be inferred only from a lack of exemplar use. Such an inference would not seem to be controversial in the four-block condition of Experiment 1. If the judgments were not directly based on activated exemplars, there would be few other sources of group knowledge beyond the development of abstract group impressions on which participants could have relied. However, in Experiment 2, there were a number of potential sources of stereotypical trait information that participants' may have used in addition to abstract group impressions. For instance, the stereotype labels may have provided access to so-called abstract exemplars pertaining to engineers and priests. Thus, participants may have based their judgments about the intelligence of engineers on the knowledge that engineers design things. Although this knowledge is still abstract in that it does not pertain to a specific group exemplar, it is not abstract *trait* knowledge (e.g., the group is intelligent). Alternatively, knowledge about stereotypical traits may have been derived from semantic associates of the category labels (priest-minister and priest-pope). Although the describes task judgment latencies and trait ratings suggest that it is unlikely that participants' judgments were based entirely on extraexperimental information (exemplar or otherwise), the possibility cannot be unequivocally discarded.

Conclusions about the mental representation of preexisting stereotypes (like those for engineers and priests) are more equivocal. Even if one accepts that participants' stereotype-consistent judgments were based on abstract group impressions in Experiment 2, the source of those abstract impressions is unclear (particularly in the one-block condition). The mixed model argues that the stereotype labels provided direct access to abstract knowledge about the stereotypic traits. This abstract knowledge presumably formed the basis for the abstract group impressions that were used to make stereotypic judgments about the group. However, abstract group impressions may have been based on a variety of different sources of information associated with the stereotype labels. As discussed above, those labels may provide access to abstract exemplars or semantic associates in addition to or instead of abstract trait knowledge. It is not unreasonable to believe that participants might have developed abstract group impressions on the basis of these abstract behaviors and semantic associates (e.g., deciding that the group of engineers was intelligent because engineers tend to design things). Thus, even if one accepts the use of abstract group impressions in the stereotypic judgment tasks, the precise representational nature of the preexisting stereotypes from which those impressions were derived is unclear.

Implicit exemplar use. One aspect of pure exemplar models that warrants mention is the notion that the exemplars that influence judgments may not be accessible to explicit recall. As such, the recall target task latencies may not have been sensitive to the implicit activation and application of exemplars in these experiments. Thus, it may be argued that all of the judgments in Experiments 1 and 2 were based on the implicit but

not explicit use of group exemplars. However, it is unclear whether the task facilitation measure would be sensitive to the implicit activation of exemplar information. The task facilitation paradigm does not measure participants' *ability* to explicitly recall an exemplar. Rather, given that explicit recall occurs, the paradigm measures how long the process takes, depending on other experimental factors. At present, whether implicit exemplar use would be exhibited by such response latencies is unknown.

More to the point, an implicit memory explanation cannot explain why exemplar retrieval is facilitated in some conditions but not others. If the measure is insensitive to implicit exemplar activation, it should be so in all conditions. At a minimum, it must be accepted that exemplars are used to a greater extent in some conditions (one block of Experiment 1 and stereotype-irrelevant traits in Experiment 2) than others. Such a position is perfectly consistent with the mixed model, which simply suggests that the relative use of exemplars depends on the presence or absence of abstract stereotypical knowledge.

Other factors in abstraction and exemplar use. The results of this research suggest that degree of experience is one factor that influences perceivers' reliance on abstractions versus exemplars for group knowledge. However, there are many other factors that influence the relative use of abstractions and exemplars. For instance, the more recently that exemplars have been activated, the more likely they are to form the basis for social judgments (Carlston & Skowronski, 1986). As described above, exemplars are less likely to form the basis for judgments when they are acquired subsequent to abstract information (Smith & Zarate, 1990). Research on categorization and individuation suggests that individuation is most likely to occur under conditions of high personal relevance and accuracy motivation (e.g., Fiske & Neuberg, 1990; Tetlock & Kim, 1987). Perhaps exemplar-based processes are also more likely to occur in such situations. It has been suggested that high relevance and motivation lead perceivers to rely on exemplar-based processes to a greater degree when making in-group judgments than when making out-group judgments, leading to the "out-group homogeneity effect" (Mackie et al., 1993; Park, Judd, & Ryan, 1991). The functional analysis of abstract knowledge outlined above suggests that conditions of limited capacity may lead to relatively abstraction-based processing strategies. In contrast, the presence of inconsistent exemplar information may lead to relatively exemplar-based judgment processes (Fiske & Neuberg, 1990). Finally, the extent to which perceivers engage in on-line encoding processes should play an important role in determining the extent to which judgments are abstraction- or exemplar-based (Hastie & Park, 1986). In the present experiments, the explicit impression formation instructions encouraged the on-line formation of abstract group impressions. However, under different instruction sets (e.g., memorization set), participants may be less likely to engage in on-line processes and form group abstractions from the stimulus exemplars. Clearly, there are a number of important factors that interact to determine the representational basis of social judgments. Future research should aim to circumscribe more precisely the conditions under which group judgments will be relatively abstraction- and exemplar-based.

Conclusion

The results of this research support a mixed model of stereotype representation. At low levels of experience, knowledge about the typical features of a novel group is exemplar-based. However, as experience with the group increases, an abstract representation of the group's typical features is formed that becomes the basis for subsequent judgments about the group. The data also suggest that abstract stereotypes of well-known groups are stored in memory and retrieved for social judgments. These results add to a growing body of evidence attesting to the generality of experience-dependent mixed models of social cognition. Similar models have shown their ability to account for the representation of knowledge about both real people (Klein & Loftus, 1993a; Klein et al., 1992) and experimentally created targets (Sherman & Klein, 1994). Furthermore, the usefulness of this model cuts across many different social targets, including the self (Klein & Loftus, 1993a; Klein et al., 1992), other individuals (Klein et al., 1992; Sherman & Klein, 1994), and now, groups of people. There appear to be certain consistencies in the mental representation of social knowledge, whether one is considering autobiographical behaviors versus semantic self-knowledge (self-schemas), biographical behaviors versus trait impressions of others, or group exemplars versus abstract stereotypes in knowledge about groups. These results make a strong argument for the continued integration of research on self-perception, impression formation, and stereotyping.

References

- Allen, R. B., & Ebbesen, E. B. (1981). Cognitive processes in person perception: Retrieval of personality trait and behavioral information. *Journal of Experimental Social Psychology*, 17, 119-141.
- Allport, G. W. (1954). *The nature of prejudice*. Cambridge, MA: Addison-Wesley.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- Bodenhausen, G. V. (1988). Stereotypic biases in social decision making and memory: Testing process models of stereotype use. *Journal of Personality and Social Psychology*, 55, 726-737.
- Bodenhausen, G. V., & Lichtenstein, M. (1987). Social stereotypes and information-processing strategies: The impact of task complexity. *Journal of Personality and Social Psychology*, 52, 871-880.
- Bodenhausen, G. V., Schwarz, N., Bless, H., & Wanke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology*, 31, 48-63.
- Bodenhausen, G. V., & Wyer, R. S., Jr. (1985). Effects of stereotypes on decision making and information-processing strategies. *Journal of Personality and Social Psychology*, 48, 267-282.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1-36). Hillsdale, NJ: Erlbaum.
- Brewer, M. B., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes and prototypes. *Journal of Personality and Social Psychology*, 41, 656-670.
- Brigham, J. C. (1971). Ethnic stereotypes. *Psychological Bulletin*, 76, 15-33.
- Brooks, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 169-211). Hillsdale, NJ: Erlbaum.
- Busmeyer, J. R., Dewey, G. I., & Medin, D. L. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 638-648.
- Buss, D. M., & Craik, K. H. (1983). The act frequency approach to personality. *Psychological Review*, 90, 105-126.
- Buss, D. M., & Craik, K. H. (1984). Acts, dispositions, and personality. In B. A. Maher & W. B. Maher (Eds.), *Progress in experimental personality research* (Vol. 13, pp. 241-301). San Diego, CA: Academic Press.
- Cantor, N., & Mischel, W. (1979). Prototypes in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12, pp. 3-52). New York: Academic Press.
- Carlston, D. E. (1980). The recall and use of traits and events in social inference processes. *Journal of Experimental Social Psychology*, 16, 303-328.
- Carlston, D. E., & Skowronski, J. J. (1986). Trait memory and behavior memory: The effects of alternative pathways on impression judgment response times. *Journal of Personality and Social Psychology*, 50, 5-13.
- Chapman, L. J., & Chapman, J. P. (1969). Illusory correlation as an obstacle to the use of valid psychodiagnostic signs. *Journal of Abnormal Psychology*, 74, 271-280.
- Cohen, C. E. (1981). Person categories and social perception: Testing some boundaries of the processing effects of prior knowledge. *Journal of Personality and Social Psychology*, 40, 441-452.
- Collins, A. M., & Quillian, M. R. (1970). Facilitating retrieval from semantic memory: The effect of repeating part of an inference. *Acta Psychologica*, 33, 304-314.
- Darley, J. M., & Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology*, 44, 20-33.
- Dreben, E. K., Fiske, S. T., & Hastie, R. (1979). The independence of evaluative and item information: Impression and recall order effects in behavior-based impression formation. *Journal of Personality and Social Psychology*, 37, 1758-1768.
- Elio, R., & Anderson, J. R. (1981). The effects of category generalizations and instance similarity on schema abstraction. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 397-417.
- Fiedler, K., Russer, S., & Gramm, K. (1993). Illusory correlations and memory performance. *Journal of Experimental Social Psychology*, 29, 111-136.
- Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 1-74). New York: Academic Press.
- Gilbert, G. M. (1951). Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology*, 46, 245-254.
- Hamilton, D. L., Dugan, P. M., & Troler, T. K. (1985). The formation of stereotypic beliefs: Further evidence for distinctiveness-based illusory correlations. *Journal of Personality and Social Psychology*, 48, 5-17.
- Hamilton, D. L., & Mackie, D. M. (1990). Specificity and generality in the nature and use of stereotypes. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 3, pp. 99-110). Hillsdale, NJ: Erlbaum.
- Hamilton, D. L., & Rose, T. L. (1980). Illusory correlation and the maintenance of stereotypic beliefs. *Journal of Personality and Social Psychology*, 39, 832-845.
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (2nd ed., Vol. 2, pp. 1-68). Hillsdale, NJ: Erlbaum.
- Hamilton, D. L., & Sherman, S. J. (1989). Illusory correlations: Im-

- plications for stereotype theory and research. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotypes and prejudice: Changing conceptions* (pp. 59–82). New York: Springer-Verlag.
- Hamilton, D. L., Sherman, S. J., & Ruvalo, C. M. (1990). Stereotype-based expectancies: Effects on information processing and social behavior. *Journal of Social Issues*, 46(2), 35–60.
- Hamilton, D. L., & Trolie, T. K. (1986). Stereotypes and stereotyping: An overview of the cognitive approach. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 127–163). Orlando, FL: Academic Press.
- Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, 93, 258–268.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411–428.
- Homa, D., Dunbar, S., & Nohre, L. (1991). Instance frequency, categorization, and the modulating effect of experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 444–458.
- Homa, D., Sterling, S., & Trepel, L. (1981). Limitation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 418–439.
- Judd, C. M., & Park, B. (1988). Out-group homogeneity: Judgments of variability at the individual and group levels. *Journal of Personality and Social Psychology*, 54, 778–788.
- Karlins, M., Coffman, T. L., & Walters, G. (1969). On the fading of social stereotypes: Studies in three generations of college students. *Journal of Personality and Social Psychology*, 13, 1–16.
- Katz, D., & Braly, K. (1933). Racial stereotypes in one hundred college students. *Journal of Abnormal and Social Psychology*, 28, 280–290.
- Keenan, J. M. (1993). An exemplar model can explain Klein and Loftus' results. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 5, pp. 69–77). Hillsdale, NJ: Erlbaum.
- Klein, S. B., & Loftus, J. (1991). [Latencies to retrieve trait abstractions and behavioral exemplars about the self and others]. Unpublished raw data.
- Klein, S. B., & Loftus, J. (1993a). Behavioral experience and trait judgments about the self. *Personality and Social Psychology Bulletin*, 19, 740–745.
- Klein, S. B., & Loftus, J. (1993b). The mental representation of trait and autobiographical knowledge about the self. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 5, pp. 1–49). Hillsdale, NJ: Erlbaum.
- Klein, S. B., Loftus, J., & Burton, H. A. (1989). Two self-reference effects: The importance of distinguishing between self-descriptiveness judgments and autobiographical retrieval in self-referent encoding. *Journal of Personality and Social Psychology*, 56, 853–865.
- Klein, S. B., Loftus, J., Trafton, J. G., & Fuhrman, R. W. (1992). Use of exemplars and abstractions in trait judgments: A model of trait knowledge about the self and others. *Journal of Personality and Social Psychology*, 63, 739–753.
- Krueger, J., & Rothbart, M. (1988). Use of categorical and individuating information in making inferences about personality. *Journal of Personality and Social Psychology*, 55, 187–195.
- Kuiper, N. A. (1981). Convergent evidence for the self as prototype: The "inverted-U RT" effect for self and other judgments. *Personality and Social Psychology Bulletin*, 35, 63–78.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.
- Kunda, Z., & Sherman-Williams, B. (1992). Stereotypes and the construal of individuating information. *Personality and Social Psychology Bulletin*, 19, 90–99.
- Lewicki, P. (1985). Nonconscious biasing effects of single instances on subsequent judgments. *Journal of Personality and Social Psychology*, 48, 563–574.
- Lingle, J. H., & Ostrom, T. M. (1979). Retrieval selectivity in memory-based impression judgments. *Journal of Personality and Social Psychology*, 37, 180–194.
- Linville, P. W., Fischer, G. W., & Salovey, P. (1989). Perceived distributions of the characteristics of in-group and out-group members: Empirical evidence and a computer simulation. *Journal of Personality and Social Psychology*, 57, 165–188.
- Lippman, W. (1922). *Public opinion*. New York: Harcourt Brace Jovanovich.
- Lord, C. G., Gilbert, D. T., & Stanley, M. A. (1982, August). *Individual self-schemas and processing information about the self*. Paper presented at the 90th Annual Convention of the American Psychological Association, Washington, DC.
- Macht, M. L., & O'Brien, E. J. (1980). Familiarity-based responding in item recognition: Evidence for the role of spreading activation. *Journal of Experimental Psychology: Human Learning and Memory*, 6, 301–318.
- Mackie, D. M., Sherman, J. W., & Worth, L. T. (1993). On-line and memory-based processes in group variability judgments. *Social Cognition*, 11, 44–69.
- Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology*, 66, 37–47.
- Malt, B. C. (1989). An on-line investigation of prototype and exemplar strategies in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 539–555.
- Manis, M., & Paskewitz, J. R. (1987). Assessing psychopathology in individuals and groups: Aggregating behavior samples to form overall impressions. *Personality and Social Psychology Bulletin*, 13, 83–94.
- McCauley, C., & Stitt, C. L. (1978). An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology*, 36, 929–940.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994a). Illusory correlation in the perception of groups: An extension of the distinctiveness-based account. *Journal of Personality and Social Psychology*, 67, 414–429.
- McConnell, A. R., Sherman, S. J., & Hamilton, D. L. (1994b). The on-line and memory-based aspects of individual and group target judgments. *Journal of Personality and Social Psychology*, 67, 173–185.
- Medin, D. L. (1988). Social categorization: Structure, processes, and purposes. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 119–126). Hillsdale, NJ: Erlbaum.
- Medin, D. L., & Wattenmaker, W. D. (1987). Category cohesiveness, theories, and cognitive archaeology. In U. Neisser (Ed.), *Concepts and conceptual development* (pp. 25–62). Cambridge, England: Cambridge University Press.
- Myers, J. L., O'Brien, E. J., Balota, D. A., & Toyofuku, M. L. (1984). Memory search without interference: The role of integration. *Cognitive Psychology*, 16, 217–242.
- Nosofsky, R. M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 87–108.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Park, B. (1986). A method for studying the development of impressions of real people. *Journal of Personality and Social Psychology*, 51, 907–917.
- Park, B. (1989). Trait attributes as on-line organizers in person impressions. In J. N. Bassili (Ed.), *On-line cognition in person perception* (pp. 39–59). Hillsdale, NJ: Erlbaum.

- Park, B., & Hastie, R. (1987). Perception of variability in category development: Instance- versus abstraction-based stereotypes. *Journal of Personality and Social Psychology*, 53, 621-635.
- Park, B., Judd, C. M., & Ryan, C. S. (1991). Social categorization and the representation of variability information. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 2, pp. 211-245). Chichester, England: Wiley.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353-363.
- Pryor, J. B. (1986). The influence of different encoding sets upon the formation of illusory correlations and group impressions. *Personality and Social Psychology Bulletin*, 12, 216-226.
- Rosch, E. R. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192-233.
- Ross, B. H., Perkins, S. J., & Tenpenny, P. L. (1990). Reminding-based category learning. *Cognitive Psychology*, 22, 460-492.
- Rothbart, M., Fulero, S., Jensen, C., Howard, J., & Birrell, B. (1978). From individual to group impressions: Availability heuristics in stereotype formation. *Journal of Experimental Social Psychology*, 14, 237-255.
- Sagar, H. A., & Schofield, J. W. (1980). Racial and behavioral cues in Black and White children's perceptions of ambiguously aggressive acts. *Journal of Personality and Social Psychology*, 39, 590-598.
- Schwarz, N., & Bless, H. (1992). Scandals and the public's trust in politicians: Assimilation and contrast effects. *Personality and Social Psychology Bulletin*, 18, 574-579.
- Sherman, J. W., & Klein, S. B. (1994). The development and representation of personality impressions. *Journal of Personality and Social Psychology*, 67, 972-983.
- Skov, R. B., & Sherman, S. J. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, 22, 93-121.
- Smith, E. R. (1990). Content and process specificity in the effects of prior experiences. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 3, pp. 1-59). Hillsdale, NJ: Erlbaum.
- Smith, E. R. (1991). Illusory correlation in a simulated exemplar-based memory. *Journal of Experimental Social Psychology*, 27, 107-123.
- Smith, E. R. (1992). The role of exemplars in social judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 107-132). Hillsdale, NJ: Erlbaum.
- Smith, E. R., & Zarate, M. A. (1990). Exemplar and prototype use in social categorization. *Social Cognition*, 8, 243-262.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99, 3-21.
- Snyder, M., & Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology*, 36, 1202-1212.
- Snyder, M., Tanke, E. D., & Berscheid, E. (1977). Social perception and interpersonal behavior: On the self-fulfilling nature of social stereotypes. *Journal of Personality and Social Psychology*, 35, 656-666.
- Snyder, M., & Uranowitz, S. W. (1978). Reconstructing the past: Some cognitive consequences of person perception. *Journal of Personality and Social Psychology*, 36, 941-950.
- Srull, T. K. (1981). Person memory: Some tests of associative storage and retrieval models. *Journal of Experimental Psychology: Human Learning and Memory*, 7, 440-462.
- Stroessner, S. J., Hamilton, D. L., & Mackie, D. M. (1992). Affect and stereotyping: The effect of induced mood on distinctiveness-based illusory correlations. *Journal of Personality and Social Psychology*, 62, 564-576.
- Taylor, S. E. (1981). A categorization approach to stereotyping. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 88-114). Hillsdale, NJ: Erlbaum.
- Taylor, S. E., & Crocker, J. (1981). Schematic bases of social information processing. In E. T. Higgins, C. P. Herman, & M. P. Zanna (Eds.), *Social cognition: The Ontario Symposium* (Vol. 1, pp. 89-134). Hillsdale, NJ: Erlbaum.
- Tetlock, P. E., & Kim, J. I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology*, 52, 700-709.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, 93, 239-257.
- Trope, Y., Cohen, O., & Alfieri, T. (1991). Behavior identification as a mediator of dispositional inference. *Journal of Personality and Social Psychology*, 61, 873-883.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology*, 47, 237-252.
- Word, C. O., Zanna, M. P., & Cooper, J. (1974). The nonverbal mediation of self-fulfilling prophecies in interracial interaction. *Journal of Experimental Social Psychology*, 10, 109-120.

Received June 19, 1995

Revision received December 13, 1995

Accepted December 13, 1995 ■