1 Stereotypes

David L. Hamilton
Jeffrey W. Sherman
University of California, Santa Barbara

Contents

What is a stereotype?
How do people develop stereotypes, and why?
Why are stereotypes so prevalent and persistent?
When and how are stereotypes used?
How can we change people's stereotypes?

These are questions that all of us—lay people and social scientists alike—have pondered when we think about stereotypes and their effects on our perceptions

of various ethnic groups. Although seemingly simple, these questions cut to the heart of the nature and functioning of stereotypes. They are questions of both societal and scientific concern and have been for a long time. Despite their importance, neither society nor science has dealt adequately with the issues and problems that these questions raise.

This chapter assesses what we have learned about some of those issues from social psychological research, and particularly from research guided by a social cognition approach to the topic. Our review and analysis is organized around questions much like those posed earlier. However, we shall see that, when approached scientifically, these broad questions actually mask a number of more specific questions, each of which presents its own issues for analysis.

A decade ago, Ashmore and Del Boca (1981) provided a useful summary of three conceptual approaches that have guided past theorizing and research on stereotyping and intergroup perceptions. They identified these approaches as the psychodynamic, sociocultural, and cognitive orientations.

The *psychodynamic* approach emphasizes the role of motivational forces and psychological benefits that can lead to and perpetuate the use of stereotypes. With its roots in Freudian thinking, this perspective includes the use of defense mechanisms such as projection and displacement of self-related sources of tension onto others, scapegoating, and an emphasis on how early childhood experiences affect intrapsychic needs in explaining intergroup perceptions. The *sociocultural* approach focuses on the variety of means by which intergroup beliefs and attitudes are acquired and maintained through social learning and social reinforcements. The focus is on how stereotypes and prejudice can be learned and perpetuated through socialization experiences, peer group influence, and media portrayals. The *cognitive* approach views stereotypes as belief systems or cognitive structures that can guide information processing, and it examines how those structures arise and how their influence on information processing affects perceptions of and interactions with members of stereotyped groups.

These orientations offer complementary, rather than competing, explanations for various phenomena involving intergroup perception. Therefore, it is likely that any single orientation is limited to providing only a partial account of these phenomena (Hamilton & Trolier, 1986; Stroebe & Insko, 1989). Nevertheless, most research on how stereotypes form and function has been guided by the concepts and methods of one of these approaches. In recent years, the cognitive approach has been particularly active and influential in generating such research. Like the other contributions to this handbook, this chapter emphasizes the advances in understanding that have been generated by an information processing analysis. However, as a means of placing this work in context, at several points we include commentary reflecting the other orientations and contrasting their emphases with those of the cognitive approach.

From the cognitive perspective, a stereotype can be defined as "a cognitive structure that contains the perceiver's knowledge, beliefs, and expectations about a human group" (Hamilton & Trolier, 1986, p. 133). Stereotypes are abstract knowledge structures linking a social group to a set of traits or behavioral characteristics. As such, stereotypes act as expectancies that guide the processing of information about the group as a whole and about particular group members (Hamilton, Sherman, & Ruvolo, 1990). In addition to these generalized expectancies, one's knowledge about particular group members (or exemplars) also may influence judgments about groups and their members.

COGNITIVE PROCESSES IN STEREOTYPE FORMATION

How do people develop stereotypes? We suspect that most people, if asked this question, would respond in one of two ways. On the one hand, they might point to the important influence of parents and other significant figures (e.g., teachers, peer groups) on the attitudes and beliefs that people develop in their formative years. On the other hand, they might cite the impact of public media, noting that people come to believe what they see about various ethnic groups as they are portrayed on television and in movies.

These would not be unreasonable answers to the question. In fact, these mechanisms are at the heart of sociocultural explanations of stereotyping. The seeds of people's conceptions of various racial and gender groups are planted in early childhood by influential adults in their lives, and they are fostered and perpetuated through their repeated perceptions of members of these groups in certain social roles as they are portrayed in the media. The role of social learning processes in the formation and maintenance of stereotypes has been a major focus of intergroup research for a long time.

In addition to these processes, the cognitive approach to stereotyping has focused on other mechanisms that can contribute to the initial formation of stereotypic belief systems. Although this approach also has a long history, it has been the catalyst for an enormous amount of research during the last 15 years. This resurgence is due, at least in part, to several empirical discoveries during the 1970s that pointed to cognitive mechanisms and biases that, in and of themselves, could contribute to the formation of stereotypes, independently of actual intergroup conflict or of social influences from significant others.

Stereotyping depends on the perception that a group of persons comprises a meaningful social entity. If individual persons were not perceived as belonging to some social unit, then there would be no basis for developing a stereotype. Inherent in this process is the perceptual separation of different social categories.

Hence, stereotyping begins with the perceptual differentiation between groups of persons. That differentiation does not, in itself, mean that a stereotype will be formed; we know, for example, that there are blue-eyed and brown-eyed people, but we do not have rich stereotypes associated with these groups. The formation of the stereotype involves an additional step of associating certain attributes

or features with those differentially perceived groups. That is, we develop beliefs about the attributes that are characteristic of each group and therefore distinguish between the groups. But these belief systems will not form if we have not already identified a collection of individuals as a group and differentiated those persons from some others. Therefore, any process that contributes to the differentiation between groups constitutes a potential basis for the formation of stereotypes. One of the key contributing factors to the resurgence of the cognitive approach to stereotypes was the fact that several lines of research provided evidence of how cognitive mechanisms can contribute to this process.

Categorization

In their perceptions of others, people often "see" others not (or not solely) as individual persons, but rather as members of social groups. Given that each person belongs to numerous social groups (based on gender, race, nationality, religion, occupation and socioeconomic status, political orientation, lifestyle, interests, etc.), viewing others in terms of such category memberships certainly captures some of the important elements of social structure and social life. In that sense, categorizing others into groups simply reflects social reality. However, beyond that, research has shown that aspects of people's cognitive mechanisms and functioning contribute to and derive from this categorization process (Hamilton & Trolier, 1986; Miller & Brewer, 1986; Oakes & Turner, 1990; Taylor, 1981). These mechanisms can have important implications for people's perceptions of and behavior toward group members.

The important role of categorization in stereotyping has been recognized for many decades. It was implicit in Lippmann's (1922) insightful analysis of the use of stereotypes in perceptions of groups. Its role was emphasized explicitly by Allport (1954) in his classic analysis of stereotyping and prejudice. But it was the simultaneous and independent development of two quite different lines of work in the early 1970s that led to an explosion of research on the role of categorization in intergroup perception. One was the development of cognitive psychology and its focus on how acquired information is organized and stored in terms of long-term "knowledge structures." The other was a program of research initiated by Tajfel (1969, 1970) that provided empirical documentation of the fundamental impact of the categorization process on social perception and behavior. Both of these initiatives focused attention on the role of cognitive factors in stereotyping and intergroup perception.

To understand the role of categorization in stereotyping, several questions need to be addressed. First, why do people categorize others into groups at all? Why not simply perceive and understand them as individuals? Second, what are the bases of social categorization? What determines the particular categorization(s) that will be used in any given situation? Third, when will the perceiver categorize others according to social groups, and when will others be perceived as

individual persons? The pervasive importance of the categorization process becomes apparent as these and other issues are discussed in this chapter. We begin with the question of why people categorize others into groups.

Why Categorize?

Why would perceivers overlook the individuality of the persons they encounter and move instead to viewing them in terms of social categories? Two major forces driving such categorization are prevalent in the literature. One emphasis views categorization as a cognitive mechanism serving the informational needs of the perceiver. Specifically, the perceiver (a) must use a limited cognitive processing system to cope with a rich and complex social stimulus environment, yet (b) needs to understand and anticipate interactions with that environment. Categorization can facilitate meeting both of those demands. The second emphasis views social categorization as deriving from people's desire to evaluate themselves positively, and therefore as motivated to see their own group as different from—and better than—other social groups. We consider each of these bases of categorization in the following subsections.

Categorization as Cognitive Efficiency. One reason to categorize others based on their apparent similarities derives from the sheer complexity of the social environment. The richness of social stimulation provided by that environment places processing demands on the human cognitive system. Therefore, attention is directed at some aspects of the social environment while others are ignored. Simplifying strategies for dealing with this information overload become functionally adaptive. As a consequence, perceivers group objects in their stimulus world into categories on the basis of their similarities and differences. Thus, categorization can be a response to the demands of information overload.

This function of perceiving others in terms of groups was demonstrated by Rothbart, Fulero, Jensen, Howard, and Birrell (1978). Subjects were presented information describing the attributes of individuals and were subsequently asked to rate the group composed of these persons on a series of attributes. The stimulus items (person-attribute pairings) were arranged such that the experimenters could determine whether subjects' ratings of the group were based on an accumulation of their conceptions of individual persons or, alternatively, on their conception of the group as a whole. When the number of stimulus items was relatively small (low memory load condition), subjects' judgments indicated that they organized the stimulus information in terms of the individual persons described. In contrast, when a large number of stimulus descriptions was presented (high memory load condition), subjects organized the descriptive information in terms of the group as a whole. Thus, under conditions of strained capacity, perceivers were less likely to develop person-based conceptions than to establish a categorical representation at the group level.

According to this cognitive view, then, categorization is, in part, a response to information overload, serving to simplify the perceiver's processing task. In this sense, categorization involves information loss. But this is not the whole story. An important consequence of categorization is that it also affords information gain. That is, through categorization, persons typically are perceived in terms of social groups about whom the perceiver, through past experiences and social learning, has developed knowledge and beliefs. This accumulated knowledge and beliefs can then be applied in understanding individual group members through inference processes. Assuming that these representations are veridical to some extent (or at least might be "functionally accurate"; Swann, 1984), these stereotype-based elaborations are likely to be useful in going beyond the information available. Thus, from the cognitive perspective, categorization involves both information loss and information gain.

Categorization as Self-Enhancement. In addition to these cognitive mechanisms that promote categorization, several motivational factors have been proposed that contribute to this same process (Maass & Schaller, 1991; Stangor & Ford, 1992; Stroebe & Insko, 1989). The most prominent theoretical account emphasizing motivational roots of categorization is social identity theory (Tajfel & Turner, 1979, 1986; see also Turner, 1987). The central hypothesis of this theory is that a person's self-esteem is, in part, derived from his or her membership in social groups. Because people typically want to maintain positive self-regard, they are motivated to hold favorable evaluations of the groups to which they belong. But there is no objective yardstick for gauging the desirability of any particular social group; such evaluations are inherently subjective. Therefore, people enhance their own group's favorability by psychologically establishing its relative superiority in comparison with some out-group. Thus, people are motivated to accentuate the evaluative difference between in-group and out-group, thereby creating intergroup discrimination based on the desire to maintain positive identity.

The development of social identity theory was stimulated by the consistent finding of in-group bias—that people evaluate their in-groups more favorably than out-groups, even when the intergroup distinction is arbitrary or based on a trivial criterion (Tajfel, 1970; see later section on in-group/out-group differentiation). Evidence for this effect is pervasive (Brewer, 1979; Messick & Mackie, 1989) and consistent with social identity theory, although the important role of self-esteem maintenance in producing these intergroup effects has not been established clearly by research findings (Maass & Schaller, 1991; Messick & Mackie, 1989).

Bases of Social Categorization

As you are on your way to work early one morning, you notice a group of joggers running through a park. One of them is an attractive, 30-ish woman, well attired in her sweatshirt and sweatpants, headband, and running shoes. In noticing

her, you might think of this woman in terms of any of several social categories to which she might belong. For example, you might categorize her as a woman, or as a runner, or more specifically, as a female runner. From her age group and expensive sportswear, you might regard her as a "yuppie." Or, in more general terms, you might think of her as "a person who tries to stay physically fit." Any of these might be plausible—and reasonably accurate—categorizations. On the other hand, you are unlikely to immediately think of her as "a resident of my community" (which also would be accurate).

This example illustrates that each person is a member of numerous social groupings, any one of which might serve as an appropriate categorical basis for perceiving the person. What determines which of several possible categorizations will be used in one's perception of this person? At this point, it is impossible to provide a definitive answer to this question, but several possibilities illustrate the range of alternative bases for categorization at the perceiver's disposal.

Primitive Categories. Some theorists (Bower & Karlin, 1974; Brewer, 1988; Bruner, 1957; Fiske & Neuberg, 1990) have suggested that a small number of critical or primitive categories are used automatically and universally in perceiving other human beings. For example, when we encounter a person it seems nearly impossible not to notice the person's gender, race, and age group, these being the three most commonly cited candidates for primitive categories.

Why would certain categories achieve this status of "primitive" categories that are employed automatically and universally? There are several highly interrelated reasons, each of which has some plausibility but has difficulties as well. First, the widespread use of these categories might simply reflect the fact that these are broad categories of human beings, which can serve as a basis for finer distinctions within these primitive categories. However, this view has difficulty explaining why certain universal features (e.g., gender, race) become the basis for primitive categorization, whereas others (e.g., eye color, hair color) do not. Also, some research suggests that once different levels of categorization or subtypes have been established, perceivers will not rely on the broader higher level classifications in categorizing others, but instead will use more specific "basic level" categories (Brewer, Dull, & Lui, 1981; Devine & Baker, 1991).

Second, primitive categories may have achieved their great importance because they have primacy status. That is, they are the first categorizations perceivers can make because they reflect features that usually are salient in a person's appearance and hence are immediately obvious to the perceiver. This explanation gives special importance to early experiences and to the first features people see in others, but neglects the importance of more recent experiences and recency effects.

Finally, perhaps the key element in this view is the argument that primitive categories are in fact the categories of greatest importance for capturing significant information about others. The idea here is that there are certain human charac-

teristics that reflect fundamental distinctions in social behavior, and it is therefore functionally useful for the perceiver to categorize others in terms of those distinctions. For example, if there are meaningful differences in social interactions with males versus females, or older versus younger persons, then recognizing another person's status on these dimensions would be useful in anticipating the nature of those interactions. The difficulty in substantiating this argument, of course, lies in documenting that these features are, in fact, most predictive of differences in social behavior.

Whether these categories have any special status as "primitive" categories or not, they will be used frequently as information about others is encoded. Frequency of category use is one of the primary determinants of that category's use in the future; use of a category increases its accessibility for future categorization (Bargh, Bond, Lombardi, & Tota, 1986; Higgins, Bargh, & Lombardi, 1985; Srull & Wyer, 1979, 1980). With repeated use across time and context, such cognitive procedures can become so routine (or "proceduralized") that they are virtually automatic (Smith, 1990).

In-group/Out-group Differentiation. Another fundamental basis for social categorization is the distinction between groups to which one belongs (in-groups) and those to which one does not (out-groups). The importance of in-group/out-group differentiations for understanding intergroup conflict and behavior has been recognized in social psychology for a long time (Allport, 1954; Campbell, 1965; Sherif, 1967). Although some writers (Allport, 1954) emphasized categorization as an important element in intergroup perception, it was only with the work of Tajfel (1969, 1970; Billig & Tajfel, 1973; Tajfel, Billig, Bundy, & Flament, 1971) that the central role of the categorization process was documented dramatically in empirical research.

Tajfel established a "minimal intergroup paradigm" (see Diehl, 1990) for studying the influence of social categorization processes, independent of actual intergroup conflict. In this paradigm, the experimental procedure occurs in two phases. In the first phase, several subjects make individual judgments about a number of stimuli (e.g., evaluative ratings of paintings by two artists; estimations of the number of dots in a stimulus display) and subsequently are given feedback that identifies them as similar to some of the other subjects and as different from the rest (e.g., in their preferences for one of the painters, or in tending to overestimate the number of dots). In actuality, subjects are randomly assigned to one or the other group, and this is done in such a way that subjects do not know which other subjects are members of their own or the other group. In this way, two groups of subjects are formed on the basis of a task that (a) seemingly might represent some "real" psychological difference yet (b) seemingly would be unrelated to performance on the tasks in the second phase of the study. In this second phase, subjects are asked to make evaluative ratings of members of their own and the other group or are asked to allocate resources to members of their

own and the other group. The primary finding from numerous studies using this paradigm is in-group bias, such that subjects allocate more resources and evaluate more favorably members of their own group, even though they do not know the specific identity (only the group membership) of the persons about whom they are responding.

Initial findings of the Tajfel group became the catalyst for subsequent research that has continued for 20 years on a variety of related topics, including intergroup differentiation, in-group bias, the perception of variability within groups, and the out-group homogeneity effect (for recent reviews, see Diehl, 1990; Messick & Mackie, 1989; Vanbeselaere, 1991). These findings also have served as the basis for several important theoretical developments in the intergroup literature (Miller & Brewer, 1986; Tajfel & Turner, 1986; Turner, 1987).

Context-Based Differentiation. One of social psychology's primary themes is the importance of the social context on people's thinking, perceptions, and behavior. It is not surprising, then, that research has shown that the nature of the social context can influence the way group members are perceived. Specifically, categorization can be based on whatever features happen to be salient within the stimulus context.

This effect was demonstrated with nonsocial stimuli by Tajfel and Wilkes (1963), who showed subjects a series of eight lines of increasing lengths. In one condition, the four shorter lines were labeled A, whereas the four longer lines were labeled B. In another condition, no labels were associated with the lines. The subjects' task was to estimate the length of each line. Relative to the no-label condition, subjects in the label condition overestimated the difference in length between the lines of the two categories (i.e., they exaggerated the difference between the fourth and fifth lines).

Similar effects have been demonstrated in the social domain. For example, Wilder (1978) showed that perceivers assume greater belief similarity between two members of the same group than of different groups, even though the basis of group membership is not diagnostic for the belief judgment being made (see also Wilder, 1981). Research by Taylor (1981; Taylor, Fiske, Etcoff, & Ruderman, 1978) showed that salient features (race, gender) can become the basis for processing and storing information about group members. For example, in one study, subjects observed a discussion among three Black and three White males and later were asked to identify which person had made certain comments during the discussion. Subjects made more within-category than between-category errors on this matching task. These findings suggest that salient race cues induced social categorization, which in turn influenced subjects' processing and representation in memory of the information about the participants.

Stimulus Salience. The immediate social context also can make a person distinctive. A person of one race or gender, for example, may be in a group otherwise composed of persons belonging to a different race or gender. Alternatively,

the social context may exist in the perceiver's head, in the form of a priori or normative expectations, such that a person becomes salient by deviating from one's stereotypic conception of that person's group. In any event, such salience, however acquired, has consequences for how that salient person is perceived.

Taylor and her colleagues (Taylor, 1981; Taylor et al., 1978; see also Lord & Saenz, 1985) conducted a series of studies demonstrating such effects. As subjects listened to an audiotape of a group discussion, a slide showing the person speaking at any given time was projected. By varying the race of the photos shown, the racial composition of the group could be manipulated experimentally while holding constant the content of the discussion. Thus, subjects in one condition saw three White and three Black males, whereas those in another condition saw five White males and one Black male. In the latter condition, the one Black person is highly salient due to group composition and, in particular, is more salient than when he appeared in the integrated group. The interesting comparison is in perceptions of the same stimulus person (same photo, same voice, same content) in the two different contexts, only one of which heightens his salience. The results showed that the salient stimulus person drew more attention, was perceived as being more prominent in the group, and was rated more extremely on trait scales than was the same person in the integrated condition.

Deviation from White Male Norm. Salience also may be created by membership in certain social groups. Zarate and Smith (1990) argued that groups that deviate from what is perceived as a cultural norm are salient. They argued that in American culture there is a White male norm. That is, it is more "normal" to be White than Black, and more "normal" to be male than female. In terms of race, this effect is created, in part, by the simple numerical preponderance of White people in the country. The historical political dominance of Whites may add to this effect. The same can be said for why males are perceived as more typical Americans than females. Thus, Blacks and women are salient to an extent because they "deviate from" this norm. Zarate and Smith argued that this salience affects categorization processes. For instance, Black males deviate from the White male norm in terms of race but not gender. Thus, Black males are more likely to be categorized as Blacks than as males, because their Blackness is salient. Similarly, White females are more likely to be categorized as females than as Whites. Their gender is salient, whereas their race is not.

Zarate and Smith (1990) tested these hypotheses using a category verification task. On each of a series of stimulus trials, subjects were presented a category label (white, black, man, woman) followed by a photograph of a Black male, White male, Black female, or White female. The subjects' task was to respond "yes" or "no," depending on whether the person shown in the photograph fit the category label. The speed of responding "yes" to a match between label and photograph was interpreted as a measure of the subject's dominant categorization of the stimulus person. For example, subjects should more quickly identify a match

between the label "black" and a photograph of a Black person than they identify a match between "white" and a photograph of a White person. Similarly, subjects should be faster in matching photographs of women with the label "female" than they are in identifying that photographs of males fit the "male" label. In general, then, the White male norm hypothesis predicts that males should be categorized more quickly by race, whereas females should be categorized more quickly by gender. In two studies reported by Zarate and Smith (1990), these hypotheses received partial support. As predicted, following race labels, subjects responded more quickly to male than to female photographs; and following gender labels, females were categorized more quickly than males. On the other hand, two general effects were problematic for the hypothesis: Overall, subjects responded more slowly to photographs of Blacks than of Whites, and more slowly to photographs of women than of men. Given that those are the cues by which persons deviate from the White male norm, these slower overall response times are difficult to understand within this framework.

Stroessner (1992) tested the White male norm hypothesis using a modification of Zarate and Smith's (1990) procedure. Rather than providing specific category labels, the cues instructed subjects to identify either the race or the gender of the person shown in the photograph. In general, Stroessner's findings were consistent with predictions. For example, judgments of Black males' race was fast and of their gender was slow, compared with those of White males (for whom response times for these two identifications were equivalent). Stroessner's results were particularly informative about the case of Black females, who differ from the perceived norm on both dimensions. When instructed to identify either the race or the gender of stimulus persons, responses to Black females were slow for both dimensions. In a second study, however, Stroessner presented compound category cues (e.g., White male, Black female, etc.) and had subjects identify (yes/no) whether the photograph matched the label. In this case, subjects were able to respond more quickly to Black females than to the other targets. Taken together, these findings support the view that Black females are spontaneously categorized on both race and gender, whereas Black males and White females are categorized spontaneously on only one dimension.

Illusory Correlation

We pointed out earlier that any process that leads to perceptual differentiation between groups constitutes a potential basis for the formation of stereotypes. The previous section discussed several ways in which categorization processes can create perceived social units, which then can become the focus of differing belief systems (i.e., stereotypes). Another process that can lead to differential group perceptions is the distinctiveness-based illusory correlation.

Hamilton and Gifford (1976) reported evidence that subjects developed differing evaluations of two groups of stimulus persons, even though the two groups were

unfamiliar (being identified only as Group A and Group B) and were described by evaluatively equivalent information. These results were due to an illusory correlation (Chapman, 1967) between group membership and behavior desirability, based on the co-occurrence of distinctive (infrequent) stimulus events. Specifically, in the information presented to the subjects, one of the groups occurred less frequently than the other, making it somewhat distinctive. Similarly, undesirable behaviors occurred less often than desirable behaviors, making their occurrence distinctive. Due to these differing relative frequencies, when a member of the smaller group performed an undesirable behavior, it constituted the co-occurrence of distinctive stimulus events. Subjects overestimated the frequency of this category of stimuli, leading them to form less favorable evaluations of the smaller group.

Hamilton and Gifford's (1976) results seemed counterintuitive, because the proportion of desirable and undesirable behaviors describing each group was the same and the normative desirability of the behaviors describing the two groups was carefully matched. Thus, due to something about the way information was utilized, subjects developed differential evaluations of groups that were objectively equivalent. As noted earlier, anything that produces differential perceptions of groups can contribute to the formation of stereotypes. Therefore, these results demonstrated that an information processing bias can create such differential perceptions, and hence can lay the groundwork for stereotype formation.

Hamilton and Gifford (1976) explained their findings as being due to the distinctiveness of the co-occurrence of infrequent stimulus information (a member of the smaller group performing an infrequently occurring type of behavior). Given this greater salience, these items become well represented in memory and are easily retrievable when judgments are subsequently called for. Hence, to the extent that subjects apply the availability heuristic (Tversky & Kahneman, 1973), these items have differential impact on those judgments. Consequently, subjects overestimate the frequency with which members of the smaller group performed undesirable acts, and therefore they make less favorable evaluations of the smaller group. These findings have stimulated a considerable amount of research investigating these illusory correlations and the mechanisms that produce them (Hamilton & Sherman, 1989; Mullen & Johnson, 1990).

The distinctiveness-based explanation advanced by Hamilton and Gifford (1976) assumes that more extensive processing of the distinctive stimuli during encoding makes them more accessible in memory and therefore more likely to be retrieved at the time judgments are made. These judgments are assumed to be at least partially based on the retrieved information (i.e., memory-based judgments). A considerable amount of research evidence is consistent with this interpretation (reviewed by Hamilton & Sherman, 1989; Mullen & Johnson, 1990). For example, the finding that the smaller group is evaluated less favorably than a larger one, even though the two groups are described by evaluatively equivalent information, is quite reliable across a number of studies (see Mullen &

Johnson, 1990, for a meta-analysis of these studies). When the infrequent behaviors are desirable, the smaller group is evaluated more favorably, rather than less favorably, indicating that the former result is not due simply to group size effects (Hamilton & Gifford, 1976). Also, subjects consistently overestimate the number of undesirable behaviors performed by the smaller group, whereas they typically are quite accurate in their estimates about the majority group. When subjects are asked to indicate the group membership of the person who performed each of the behaviors, they tend to overattribute the undesirable behaviors to the smaller group. When asked to recall the behavioral information, subjects remember a higher proportion of the distinctive behaviors (undesirable behaviors performed by the smaller group; Hamilton, Dugan, & Trolier, 1985). All of these findings are consistent with the argument that the co-occurrence of infrequent stimulus categories has special impact on subsequent judgments.

Other results suggest that illusory correlations reflect judgments that are memory-based, being influenced by the greater accessibility of the distinctive items in memory. For example, the typical illusory correlation results just summarized do not occur when the stimulus behaviors describe individuals, rather than members of groups (Sanbonmatsu, Sherman, & Hamilton, 1987), presumably because processing information about individual persons invokes on-line integrative processing so that later judgments are not based on retrieval of accessible items. Consistent with this view, Pryor (1986) showed that illusory correlations of groups do not form when subjects are instructed to develop well-integrated impressions of the groups—instructions that also would induce on-line processing. In a similar vein, several studies have shown that self-relevant motives can induce on-line processing and thereby modify the nature of illusory correlation results (Maass & Schaller, 1991; Sanbonmatsu, Shavitt, Sherman, & Roskos-Ewoldson, 1987; Schaller, 1991; Schaller & Maass, 1989).

Although these findings support the distinctiveness interpretation (Hamilton & Sherman, 1989), two alternative explanations for these illusory correlation effects have been proposed recently, neither of which posits any special processing of the infrequently occurring items. Smith (1991) showed that a computer simulation based on Hintzman's (1986) memory model was able to produce results similar to those obtained in illusory correlation studies. This model does not assume any special attentional, encoding, or retrieval processes associated with distinctive stimuli. Rather, judgments are based on activated memory traces of previously learned stimulus items. Because of their differential frequencies of occurrence, the difference between the number of desirable and undesirable behaviors performed by Group A is greater than that difference for Group B. Therefore, comparable rates of retrieval for all four categories of items would lead this difference to be greater for Group A than Group B. If evaluative judgments reflect the difference between those differences, then the typical illusory correlation effects would be obtained. In fact, this is what Smith's (1991) computer simulation has shown.

A second explanation that does not rely on differential encoding of distinctive information (Fiedler, 1991) posits that, because of the differing frequencies of the various stimulus categories and imperfect learning of those frequencies, there is "information loss" that produces "regression" in subjects' estimates of those frequencies. Hence, subjects overestimate the frequency of the least frequent item type simply due to regression effects in judgment, rather than to any distinctive properties of those items. Other judgments (e.g., evaluative ratings of the groups) are based on those estimated frequencies, producing the typical illusory correlation results.

Although different from each other, both the Smith (1991) and the Fiedler (1991) interpretations were able to account for many (although not all) of the findings that typically had been viewed as supporting the distinctiveness interpretation. These papers point to mechanisms that, at minimum, may contribute to the occurrence of illusory correlations. The more important and challenging theoretical question is whether these reinterpretations vitiate the need to assume that the distinctive items receive any differential processing as the information is encoded. To address this question, two issues become of crucial concern: Are distinctive items processed differently than other item types? Subsequently, when judgments are made, are distinctive items more accessible than other item types?

Two recent papers provide evidence relevant to these questions. First, Stroessner, Hamilton, and Mackie (1992) measured subjects' processing times as they read and encoded the items of stimulus information. They showed that, under standard conditions, subjects spent a longer time processing the distinctive items than the other items. Thus, it appears that these items get some kind of additional processing. Moreover, Stroessner et al. showed that induced mood states disrupt the formation of illusory correlations and that they do so by undermining the differential attention directed to these distinctive items. This research provides the most direct evidence for the differential processing hypothesized by the distinctiveness interpretation. Second, Johnson and Mullen (in press) measured subjects' response times in judging whether each stimulus behavior had been performed by a member of the larger or the smaller group. They found that subjects responded more quickly to the undesirable behaviors performed by the smaller group, indicating that these behaviors were more accessible in memory than the other three types of behaviors.

These findings (Johnson & Mullen, in press; Stroessner et al., 1992) provide direct evidence for the differential processing and, subsequently, the increased accessibility of distinctive information. Therefore, they provide strong support for the distinctiveness interpretation. In addition, however, recent analyses have shown that other mechanisms also may contribute to the formation of illusory correlations under these conditions (Fiedler, 1991; Smith, 1991). An important task for future research will be to determine the conditions under which each of these mechanisms is likely to be influential.

STEREOTYPES AS COGNITIVE STRUCTURES

Up to this point, we have attempted to develop the argument that stereotypes are based on the perceived differentiation between social groups, and that both categorization processes and illusory correlations can contribute significantly to that differentiation. This differentiation does not, in itself, constitute stereotyping, but it does lay the foundation for the development of stereotypes. As the perceiver acquires knowledge and beliefs about a group, and those beliefs become associated with that group, a stereotype of that group becomes established. This stereotype is stored in memory as a cognitive structure and can then influence subsequent perceptions of and behaviors toward that group and its members. Although this statement may seem intuitively obvious, it masks a host of questions about (a) the nature of that cognitive structure and (b) how it produces effects on judgments and behaviors. In this section, we discuss a variety of ways that stereotypes, as cognitive structures, have been conceptualized. The next section examines the ways that stereotypes can influence information processing, judgments, and behavior.

Early Approaches

For many years, social psychologists simply defined stereotypes as belief systems about social groups and devoted their efforts primarily to identifying and measuring the content of those belief systems (Brigham, 1971; Hamilton, Stroessner, & Driscoll, in press). With the resurgence of the cognitive perspective came more focused concerns about what these structures look like, how they are organized, and how knowledge and beliefs about social groups are stored in memory and subsequently utilized.

Initial advances in this direction reflected the implicit personality theory (IPT) approach to social perception. This work attempted to study the structure underlying stereotypic belief systems, identifying the complexity and dimensional structure of the stereotype by applying techniques such as multidimensional scaling and hierarchical clustering analysis (Ashmore, 1981; Ashmore & Del Boca, 1979; Jones & Ashmore, 1973). Although these analyses moved us beyond simple conceptions of a stereotype as a collection of traits associated with a group concept, this approach soon faced the same limitations that ultimately constrained the usefulness of the IPT approach. That is, the research could define and measure the structural organization of the stereotype (i.e., its content, its dimensionality, etc.), but the methods employed were not informative about how that structure influenced one's perceptions of the group, how it biased subsequent information processing, and so on.

Cognitive Representations of Groups and Group Members

To further understand these issues, recent research on stereotypes has investigated how one's knowledge about groups and group members is represented in memory and how that knowledge is used in making judgments. If stereotypes

are cognitive structures, what do those structures look like? Some approaches to this question view stereotypes as generalized conceptions of the prominent attributes of a group, and, as such, stereotypes are represented in memory as abstractions based on previous learning and experience. Other approaches emphasize the role of exemplar information, with less (if any) role afforded to abstract concepts. To investigate these representations, researchers have drawn on and used cognitive models of categorization and judgment, which concern how prior knowledge is used in categorizing and judging individual instances that may or may not be category members. By understanding these processes, we can gain knowledge about the nature of the underlying cognitive representation.

In this section, we summarize both abstraction-based and exemplar-based models of representation and their application to social categorization and judgment processes in the domain of stereotyping. We hope to demonstrate that both types of representation are necessary elements for an adequate understanding of stereotypes.

A *Abstraction-Based Representations*

Stereotypes traditionally have been conceptualized as abstract generalized beliefs about social groups, with little, if any, attention to the role of one's knowledge about individual group members, or exemplars. In such abstraction-based viewpoints, a generalized conception of a group develops as information about that group is acquired. This information may be acquired from a variety of sources, including through first-hand experience with group members and through social learning from other sources, such as family, friends, and the media. As increasing numbers of individual group members are encountered, and as more is learned about the group from other sources, an abstract group representation summarizing this information is formed and stored for future use (Hayes-Roth & Hayes-Roth, 1977; Posner & Keele, 1968). A number of different terms have been used to refer to these abstract representations, including *schemata*, *prototypes*, and *knowledge structures*. Although distinctions among these terms may be relevant in some contexts (Fiske & Taylor, 1991), for our purposes they can be used interchangeably. In this view, stereotypes are abstract cognitive representations that summarize one's knowledge about groups in generalized form. That is, they are group prototypes.

Any group judgment process must be preceded by an initial act of social categorization. That is, before a group stereotype can be activated and used in one's perceptions of another person, that target person must be categorized as a member of the group. This process can be influenced by the stereotype.

According to "prototype" models of social categorization, categories do not have defining features or criteria that determine whether a person is a member of the category. Rather, categories are "fuzzy sets" whose members vary in the degree to which they fit within the category. Some instances are better examples

of the category than others. This "fit" is determined by a comparison of the individual instance to a category prototype that summarizes the features of category members. If the features of the individual are sufficiently similar to the features of the group stereotype, then the person will be categorized as a member of that group. An individual person can be categorized in terms of several possible categories to which the person belongs. Generally, the target will be categorized according to the group prototype that provides the closest match to the target's features (Brewer et al., 1981; Cantor & Mischel, 1979; Rosch, 1978).

As noted earlier, some writers (Brewer, 1988; Fiske & Neuberg, 1990; Zarate & Smith, 1990) have suggested that categorization is an automatic process based on certain primitive features (e.g., skin color, gender, age) that serve as important distinctions and have been used frequently in the past. In this case, the "similarity matching" process involved in categorization relies on a single salient cue, rather than on resemblance to the group stereotype. Then, upon categorization, the stereotype will be activated. This argument suggests that it may be possible to activate a stereotype without having had to use it for categorization. Put another way, a stereotype may be activated and used in subsequent information processing even if it was not used for initial categorization. The relationships among categorization processes, stereotype activation, and stereotype utilization represent important theoretical and empirical issues that recur at several points in this chapter.

In summary, abstraction-based models of representation posit that stimuli are categorized by comparison to group prototypes or stereotypes. Stereotype activation permits the perceiver to infer a target's features through group membership and can bias subsequent information processing in stereotype confirming ways.

B *Limitations of Abstraction-Based Models*

The basic assumptions of abstraction-based representations have been at least implicit in the literature on stereotypes for several decades. Despite their prevalence, these models face a number of problems that limit their adequacy as a full accounting of stereotypes. Several of these constraints are summarized briefly in this subsection.

Perceivers Possess Group Variability Information. Abstraction-based models of group representation have difficulty accounting for perceivers' knowledge of and sensitivity to variability in groups. That is, abstraction-based models represent the central tendencies in the perceiver's beliefs about a target group, but fail to represent one's perception of the diversity within that group. Yet perceivers readily can judge the degree to which category members vary around the mean on different dimensions. Therefore, it would seem that to judge group variability, perceivers would need to rely on the retrieval of specific category exemplars.

However, this issue is not as straightforward as it might seem. If one considers variability within a group to be merely another feature of the group, it is entirely possible that the same processes that govern the acquisition of abstract information about a group's attributes also could create abstract knowledge about a group's variability. That is, perceivers could acquire knowledge about group variability in the same way they acquire knowledge about other features (through social learning, through inferences based on observations, etc.). If so, then variability information could be stored as part of the group stereotype along with information about other features. This issue has been the subject of considerable debate, with some researchers arguing that variability information is not stored as abstract knowledge (Linville, Fischer, & Salovey, 1989), and others arguing that it is (Park & Hastie, 1987). We return to the issue of perceived variability in a later section.

Subtypes. Much of the literature on stereotypes has assumed that perceivers hold and use very generalized conceptions of large social groups (e.g., Blacks, women, Germans, etc.). A potential limitation of such generalized conceptions is that they are overinclusive (i.e., they fail to recognize, at the stereotype level, certain distinctions that perceivers quite obviously make within those broad human categories). Is a generalized stereotype of women really viable when common knowledge indicates that this stereotype would not apply equally well to a businesswoman, a barmaid, a homemaker, the checkout clerk at the grocery store, and a female college professor?

In recognition of this problem, recent research (Ashmore, 1981; Brewer, 1988; Brewer et al., 1981; Devine & Baker, 1991; also see Deaux, Winton, Crowley, & Lewis, 1985) has shown that such broad stereotypes are often the highest or superordinate categories in a hierarchical category system that also contains more distinct, subordinate levels or subtypes. In fact, it is argued that target persons are typically categorized, and hence stereotypes are created, stored, and used at this subtype level.

Brewer et al. (1981) provided a particularly impressive demonstration of stereotypes that exist at the subtype, as well as the superordinate, level. Subjects sorted photographs of elderly persons into categories on the basis of perceived similarity in personality. There was a substantial degree of consensus among subjects that most stimulus persons fit into three different subcategories. Brewer et al. presented a new sample of subjects with sets of three photographs from the various subtypes, and the subjects' task was to indicate (on an adjective checklist) those attributes that all three of the persons possessed. An attribute was considered stereotypic of the subtype if at least 50% of the subjects endorsed the attribute as descriptive of all three persons. For each subtype, there were several attributes for which such consensus occurred, and the attribute sets characteristic of each subtype were quite distinct from each other. On the other hand, when

subjects were shown three photographs consisting of one person from each subtype (thereby representing the superordinate category), there were very few attributes that subjects agreed described all three of them. These findings indicate that, although people may at times categorize and describe others simply as "old people," they in fact possess much richer and better differentiated stereotypes of more specific subgroups of that inclusive category.

Despite such findings, Smith (1990) questioned the necessity of assuming that stereotypes form at the subtype level. He pointed out there could be infinite numbers of subtypes for any category (Black Baptist preacher, Black politician, Black college student, etc.), and many of these subtypes could be context specific (Black running back in Big Ten football games). Smith doubted that perceivers possibly could have stored distinct stereotypes for all of the possible subtypes that they might have developed, and instead argued that many subtype characterizations probably are computed based on the retrieval of exemplars at the time judgments are made. However, Hamilton and Mackie (1990) questioned the necessity of this position. Given extensive (direct or indirect) experience with a group of people, it certainly seems possible that the group representation can become highly differentiated, creating an extremely large number of specific subtypes. Even so, Smith was correct to point out that people can make judgments about groups of people that they had not considered previously as meaningful subgroups (perhaps Black running backs in Big Ten games). Because these subgroups have no existing representation, judgments about them necessarily are based on the retrieval of exemplars.

The Impact of Instances. Finally, abstraction models have difficulty accounting for the large impact that particular instances can have on processing. For example, in one study, subjects expected a newly encountered individual with short hair to be unfriendly, simply because of a previous encounter with an unfriendly individual who also happened to have short hair (Lewicki, 1985). Upon encountering the second target, subjects retrieved the first individual as a basis of judgment.

Although this study is often cited as demonstrating the importance of exemplar representations for social judgments, it may not provide strong evidence for nonreliance on abstract knowledge representations, such as stereotypes, because subjects may not have possessed a stereotype relevant to the target in question (i.e., a stereotype regarding the friendliness of long- and short-haired people). Consistent with this view, Smith and Zarate (1990) demonstrated that specific exemplars are utilized as a basis for categorization when abstract knowledge is based on few instances and is weakly defined.

Summary. Despite the historical preeminence of generalized stereotypic conceptions, there are good reasons to believe that group categorization and judgment processes do not rely solely on abstract representations of groups. Variability

judgments, the role of subtypes, and the impact of specific instances in category judgments suggest that, in addition to generalized stereotypes, people store and use knowledge of individual group members. Next, we turn to a consideration of specific models of exemplar-based processes.

C. Exemplar-Based Representations

Most exemplar-based models were developed to account for categorization processes (Brooks, 1978; Medin & Schaffer, 1978). According to these models, categorization is not achieved by comparing the target to a category prototype. Rather, categorization is achieved by comparing the target to the category membership of the set of retrieved exemplars. The exemplars retrieved are assumed to be those that are most similar to the target. Exemplar retrieval and usage need not be accessible to consciousness, and an extremely large number of exemplars may be retrieved (Medin & Schaffer, 1978; Nosofsky, 1987; Smith & Zarate, 1990). Because exemplar retrieval is often an implicit process, it may not be revealed by such typical dependent measures as recall and recognition.

In their extreme form, exemplar models suggest that stimuli are reacted to on the basis of retrieved exemplars alone, and that categorization per se does not occur. In fact, some models suggest that there is no such thing as abstract, categorical knowledge (Hintzman, 1986).

Recently, Smith has extended his exemplar-based model of social categorization to cover social judgment processes (Smith, 1990; Smith & Zarate, 1992). According to this extension, not only is exemplar retrieval responsible for the initial categorization of a stimulus, but the retrieved exemplars also guide subsequent judgment processes involving the target (Andersen & Cole, 1990; Gilovich, 1981; Lewicki, 1986). Smith (1990) proposed that stereotypes be reconceptualized as summarized exemplars, and not as stored abstract knowledge. Presumably, the same exemplars that are retrieved for purposes of categorization are utilized to form a group stereotype (Kahneman & Miller, 1986). The features of those exemplars are summarized, creating a stereotype of the particular group into which the target has been categorized, which then acts as an expectancy that can guide subsequent processes.

However, to term such a post-hoc group summary a *stereotype* is to alter the meaning of the term, which has always referred to some form of generalized representation. Moreover, if exemplars are retrieved and summarized for making group judgments, the necessity for postulating the creation of a stereotype at all is unclear. Smith's (1990) conceptualization essentially denies the existence of stored abstract knowledge. However, just as models that deny the existence of exemplar representations have their limitations, so do models that deny the existence of abstract knowledge such as stereotypes.

D. Limitations of Exemplar-Based Models

Social Learning. One basic problem for pure exemplar models is that people often acquire abstract information about a group from other sources (e.g., family, friends, television). People may be told that "women are emotional" or that "men are obstinate," rather than abstracting that knowledge from the retrieval of exemplars.

Some exemplar models account for such knowledge by suggesting that abstract information is merely stored and activated as another exemplar when categorization or judgments occur (Hintzman, 1986; Linville et al., 1989; Smith & Zarate, 1990). Thus, when a female target is encountered, perceivers may retrieve the exemplar that "women are emotional" as well as other particular instances of women. Although this conceptualization provides a mechanism for representing abstract characterizations in exemplar form, it blurs the distinction between abstract and specific information. When such an abstraction is retrieved and used in judgment, how does one know whether it was acquired and stored as a specific instance (as in this example) or is a generalization resulting from an abstraction process? This conceptualization lends credence to the view that, as abstraction-based and exemplar-based theories become more richly developed, they approach being indistinguishable from one another (Barsalou, 1990).

Exemplars Must Be Held Together by a Theory. To form a coherent group representation, exemplars must be joined together by some sort of category definition or theory (Medin & Wattenmaker, 1987). There must be some criteria for category inclusion. Smith (1990) noted that, without a category theory, a set of exemplars (e.g., a golden retriever, the number 39, and the graphics board of a computer) does not hang together and form a coherent concept or category.

For instance, if someone is asked to describe the typical woman, the exemplars retrieved must be constrained by a theory of inclusion criterion. Otherwise, when asked to make category judgments, perceivers could not activate appropriate exemplars. This point is particularly relevant in situations where there is no stimulus other than the category label to act as a retrieval cue. In this case, the category knowledge activated will serve to constrain the set of exemplars that may be retrieved (Kahneman & Miller, 1986; Rothbart & John, 1985).

In situations where an actual stimulus target is being categorized or judged, the possibility remains that the target will cue exemplar retrieval and be reacted to without any mediating influence of category knowledge. Note, however, that this view directly contradicts theories of primitive categorization (Bower & Karlin, 1974; Brewer, 1988; Bruner, 1957; Fiske & Neuberg, 1990). Questions regarding the automaticity of social categorization and stereotype activation become of crucial importance for this issue, and are discussed in a later section.

Efficiency. Exemplar-based models posit that stereotypic judgments are based on the set of exemplars retrieved at the time the judgment is made, and not on any generalized conception of the target person or group. However, after extensive processing and representation of numerous exemplars (group members), it would seem efficient (even if not necessary) at some point to generalize, to note commonalties among exemplars, and to summarize some of the attained knowledge in the form of group-level characterizations (Hamilton & Mackie, 1990). The alternative—continuing to process, record, and keep track of an ever-increasing number of group members—seems the less efficient, more demanding system.

Exemplar models argue that exemplar retrieval is a parallel and implicit process that does not require significant capacity (Nosofsky, 1987; Smith, 1990), so "keeping track" requires neither time nor cognitive resources. Nevertheless, some research has shown that cognitive load can influence the nature of category representations (Rothbart et al., 1978).

One troubling aspect of many exemplar models is that they are very difficult to disconfirm. If exemplar retrieval and summation is a parallel, implicit, and unconscious process, it is not clear how one would demonstrate that exemplars have not been retrieved for use. Typical memory measures such as recognition, recall, and response time would not necessarily provide informative data.

On-Line Processing. Other analyses indicate that abstraction does seem to occur spontaneously under many circumstances (Fried & Holyoak, 1984; Park & Hastie, 1987; Posner & Keele, 1968). The distinction between on-line and memory-based processes is central to this issue (Hastie & Park, 1986). When perceivers have the explicit goal of forming a group-level judgment as they process information about individual group members, they abstract information from the exemplars and store it (Park & Hastie, 1987). When a group judgment is required, subjects simply retrieve the stored abstraction. However, if subjects have no such processing goal at the time information is encoded, an abstraction may not be formed on-line. In this case, subsequent (memory-based) judgments about the group would require the retrieval and summarization of specific instances.

If judgments of a group are memory-based (as posited by exemplar models), one would expect a strong relationship between the exemplar information retrieved and the resulting group judgments. However, some research (Judd & Park, 1988; Park & Hastie, 1987) has found that group judgments can be independent of exemplar retrieval, suggesting that group-level characterizations had been abstracted on-line and used in subsequent judgments.

Abstraction Use in Judgments. Regardless of whether a group judgment has been created on-line or through exemplar retrieval, it presumably would be functional to store that judgment for future use. Particularly in cases in which the judgment is likely to be required again and again, it seems highly inefficient to continually recalculate the judgment through exemplar retrieval. Research has

demonstrated that pre-formed judgments, and not the original exemplars from which those judgments were derived, are retrieved for subsequent judgment processes (Carlston, 1980; Fazio, Lenn, & Effrein, 1984; Lingle & Ostrom, 1979).

Summary. Just as pure abstraction models have difficulties accounting for some findings, so do pure exemplar models. Social learning processes, the necessity of theories to make sense of exemplars, efficiency problems, on-line formation of group judgments, and the use of pre-formed abstractions in subsequent processing, all pose problems for pure exemplar models.

Mixed Models of Group Representation

Given the limitations of both pure abstraction-based and pure exemplar-based models, many researchers have concluded that a viable model of social categorization and judgment processes must include both abstract knowledge representations, such as stereotypes, and specific group exemplars (Hamilton & Mackie, 1990; Homa, Dunbar, & Nohre, 1991; Kahneman & Miller, 1986; Klein, Loftus, Trafton, & Fuhrman, 1992; Malt, 1989; Murphy & Medin, 1985). Given such a mixed model, a number of important questions arise. One obvious issue is to determine the conditions that lead to relatively exemplar-based or abstraction-based processing. Several relevant considerations can be cited.

Learning Order of Information. Based on research by Medin, Altom, and Murphy (1984), Smith and Zarate (1990) designed an experiment to examine the use of abstract and exemplar information in social categorization. The subjects' task was to learn to classify members of two different groups. Some subjects learned prototype information before they learned the group exemplars; other subjects were given no information about the group prototype. After the learning phase of the experiment, subjects were asked to categorize the learned targets as well as a set of new test targets. The stimuli were designed such that patterns of categorization would differ depending on whether subjects used an exemplar-based or abstraction-based categorization strategy. The results showed that when subjects were given no group prototype information, they categorized new exemplars by comparing them to previously learned exemplars. However, subjects who first learned group prototype information based their categorizations on comparison with that prototype.

These data suggest that the order in which perceivers learn abstract and exemplar information about a group can affect subsequent processing of group members. Thus, if perceivers possess a previously formed stereotype when they encounter group members, the stereotype may direct subsequent processing. However, when perceivers interact with members of a group about whom they have no stereotype, retrieved category exemplars may play a particularly important role in subsequent processing.

b. On-Line Versus Memory-Based Processing. As noted previously, the extent to which abstract representations are formed as new exemplars are encountered is an important determinant of the kind of information that will be used for subsequent processing. If a group-level characterization has been induced during exemplar encoding, that abstraction should be particularly likely to influence subsequent judgments (Fried & Holyoak, 1984; Hastie & Park, 1986; Homa et al., 1991; Posner & Keele, 1968). However, if no abstraction has been formed on-line, perceivers necessarily will have to rely on exemplar retrieval to guide subsequent processing. An important question for future research concerns the extent to which stereotypes, in the absence of on-line judgment formations, provide abstract information that precludes the retrieval of exemplars for judgment purposes.

c. Conditions of Limited Capacity. Cognitive capacity affects the processing of information about group members. As already noted, Rothbart et al. (1978) showed that, under conditions of constrained capacity, perceivers were more likely to organize information in terms of a group than individual members. Other research indicates that cognitive demands affect the likelihood that perceivers will rely on category-level knowledge, rather than individuating target information, when judging a stimulus person (Bechtold, Naccarato, & Zanna, 1986; Bodenhausen & Lichtenstein, 1987; Fiske & Pavelchak, 1986). There also is evidence that stereotypes are spontaneously or automatically activated (Devine, 1989; Macrae, Milne, & Bodenhausen, in press) or, upon activation, automatically applied (Gilbert & Hixon, 1991). These findings suggest that the use of abstract information requires little capacity, and that perceivers are more likely to rely on abstract information under capacity constraints.

Exemplar-based models propose that issues of cognitive capacity are irrelevant to exemplar retrieval and summation (Smith, 1990), because exemplar processing is assumed to be implicit and parallel. Given these assumptions, evidence that capacity limitations can affect exemplar use would be difficult to obtain. Clearly, the effects of capacity load on the differential use of exemplar and abstract knowledge are undetermined at this point.

d. Recent Encounters. As Lewicki (1986) demonstrated, recent encounters can affect subsequent judgments. The strength of these effects is determined by the recency of the encounter and the availability in memory of the exemplar (Kahneman & Tversky, 1973). In addition, the exemplar is likely only to affect judgments of targets that bear some similarity to the exemplar (Kahneman & Miller, 1986; Smith, 1990). This type of analogical processing likely decreases the effects of abstract category knowledge.

e. Summary. A number of factors appear to mediate the extent to which categorization and judgment processes rely on abstract or exemplar information. Whether exemplar or abstract information is learned first affects which type of information will dominate subsequent processing. Judgments that are formed on-line are unlikely to rely on exemplar retrieval, whereas memory-based judgments rely on exemplar retrieval. Under conditions of limited capacity, subjects may be less likely to rely on exemplar retrieval, and recent encounters can influence the use of abstract and exemplar information.

(g) Representation of Group Variability Information

One of the cardinal features of stereotypes, dating back to the earliest writing on the topic (Lippmann, 1922), is that stereotyping represents an overgeneralization in that attributes are ascribed to all group members, leading to perceptions that "they are all alike." However, it is also clear that perceivers are at least somewhat sensitive to differences among the members of groups. Members of some groups do seem to be all alike or homogeneous, whereas members of other groups appear to be more diverse or heterogeneous. In more general terms, groups differ in their perceived variability. This means that perceivers somehow can assess the degree of variability in those groups. How is this achieved? Is that perception of variability a part of one's stereotype of a group?

Research on perceptions of group variability was stimulated by the highly replicable out-group homogeneity effect (OHE)—the tendency of people to perceive out-groups as being more homogeneous than groups to which they belong (for discussions of this literature, see Brewer, 1993; Mullen & Hu, 1989; Ostrom & Sedikides, 1992; Quattrone, 1986). Although the OHE was the catalyst, the mechanisms underlying perceptions of group variability are also of central importance for understanding stereotypes as mental representations of groups. Perceived group variability has been shown to affect many aspects of stereotyping and intergroup perception, including (a) the probability that a stereotype will be applied to a particular group member (Park & Hastie, 1987; Park, Judd, & Ryan, 1991), (b) the likelihood of generalization from one group member to the whole group (Park & Hastie, 1987), and (c) the likelihood of generalization from one group member to another (Quattrone & Jones, 1980). Perceived variability affects both social categorization and social judgment processes (Park & Hastie, 1987), and it may have a significant impact on stereotype change (Park et al., 1991). Several of these points are addressed in other sections of this chapter. We focus here on how information pertaining to group variability is represented in memory and how variability judgments are made.

a. *Models of Group Variability Representation*

Several models have been proposed that offer differing positions on these issues. Some models argue that stereotypes, as group-level representations, cannot account for perceivers' sensitivity to the degree of variability among a group's members. In these accounts, variability information is not represented at the group

level, but is derived from exemplar retrieval. Other models hold that variability information is stored at the group level, arguing that variability information is abstracted in the same manner as information about any other attribute and is stored as part of the group stereotype. We summarize four models of perceived variability and discuss their implications for stereotype representation.

a. *The PDIST Model.* As mentioned earlier, one of the reasons some researchers have turned to exemplar models of social cognition is to account for perceivers' sensitivity to variability within groups of people (Linville et al., 1989; Smith, 1990). Exemplar models argue that group-level information cannot account for variability knowledge. Rather, variability estimates are based on retrieved group exemplars.

One of the most clearly specified models of variability representation is the exemplar model proposed by Linville and her colleagues, called PDIST (Linville et al., 1989; Linville, Salovey, & Fischer, 1986). According to the model, judgments of variability are created by retrieving particular group exemplars and summarizing their features.

In PDIST, the degree of perceived variability depends primarily on the number of group members known. If one retrieves a large number of exemplars, those instances are likely to have more variability than when a small number of exemplars is retrieved. One consequence is that typically the perceived variability of familiar groups (with many easily retrievable exemplars) will be higher than the perceived variability of unfamiliar groups (with few easily accessible exemplars). According to PDIST, the OHE is a result of differential familiarity with in-groups and out-groups: People simply know more in-group than out-group members, resulting in greater perceived variability in the in-group.

The Linville et al. (1989) model does allow for the representation of abstract knowledge (e.g., the statement that "men are obstinate"). However, that "abstract" information about the group is stored, retrieved, and weighted as just another exemplar.

In this model, then, group variability information is not stored as part of a group representation. Variability estimates are strictly retrieval based and are calculated only when they are specifically requested.

What evidence supports the PDIST model? Computer simulations have demonstrated that, in a retrieval-based judgment process, group familiarity does affect perceived variability, as predicted (Linville et al., 1989). In addition, Linville et al. (1986, 1989) showed that perceived variability increases as familiarity with groups increases over time, and that variability estimates are equal for groups that are known equally well (males and females).

However, the relationship between familiarity and perceived variability is not clear cut. Some researchers have found that greater familiarity does not lead to greater perceived variability (Jones, Wood, & Quattrone, 1981). Others have found in-group/out-group differences in perceived variability between genders, a case

where degree of familiarity should be comparable (Park & Rothbart, 1982). Still others have found in-group/out-group differences using a minimal group situation that equates knowledge of the two groups (Judd & Park, 1988; Mackie, Sherman, & Worth, 1993). Finally, Karasawa and Brewer (1992) demonstrated that in some cases greater familiarity leads to lower perceived variability. Clearly, the relationship between familiarity and perceived variability is not as direct as posited by the PDIST model.

b. *Dual Predictor Model.* A second exemplar model of variability representation is the dual predictor model recently proposed by Kashima and Kashima (1993). Like PDIST, this model assumes that variability information is not stored with the group stereotype, but is instead derived from the retrieval of particular group members. However, the exemplar summarization process in the dual predictor model is markedly different from the one in PDIST.

The dual predictor model assumes that the similarity of two exemplars is determined dually by the number of features the two exemplars share and the number of distinct features possessed by each exemplar (Tversky, 1977). As the number of shared features increases, similarity increases; as the number of distinct features increases, similarity decreases. Group variability judgments are essentially judgments of the similarity among group members. When a group variability judgment is called for, group exemplars are retrieved and their similarities and differences are assessed. The group variability estimate is based on the overall numbers of similarities and differences that occur in the retrieved exemplars.

Unlike PDIST, variability is not determined by the raw number of exemplars retrieved, but rather by the particular qualities of those exemplars. If all the exemplars retrieved are very similar, the variability estimate will be low, even if there are many of them (Quattrone, 1986). On the other hand, variability judgments may be quite high even if only a few (but highly dissimilar) exemplars are retrieved.

In support of their argument, Kashima and Kashima (1993) found that increased similarity information resulted in lower perceived variability, whereas increased difference information resulted in higher perceived variability, as expected. However, they also found a significant main effect due to familiarity: greater familiarity was associated with greater perceived variability, independently of similarity and difference information. Clearly, further research on this model is necessary.

In discussing PDIST and the dual predictor model, we have examined the effects of familiarity and similarity on perceived variability, because of their relevance to these models. However, these data do not directly assess whether variability estimation actually relies on exemplar retrieval, as these models argue. In fact, there is a paucity of evidence bearing directly on this central question. The Linville et al. (1989) computer simulation demonstrated that such a

retrieval-based process is sufficient to produce the predicted effects. However, variability estimates may not necessarily be retrieval based.

One indication of retrieval-based processes is a relationship between judgments and recalled information. If perceivers retrieve exemplar information to make their judgments, then that information should be positively related to the resulting judgment. Some research has failed to find such a relationship. Park and Hastie (1987) found that repeating certain behaviors (presumably making them more memorable) did not affect subsequent variability judgments. Judd and Park (1988) also failed to find a relationship between recalled information and variability judgments. However, the conclusion from these findings that the variability judgments were not retrieval based may be premature. For example, in the Judd and Park (1988) study, the information presented to subjects was visible when the judgments were requested (see Linville et al., 1989), therefore judgments did not need to rely on retrieval at all.

In contrast, Mackie et al. (1993) found a relationship between retrieved information and variability judgments, suggesting a retrieval-based process. Specifically, they found a relationship between recalled similarity information and perceived variability, as predicted by the dual predictor model, although not between recalled difference information and perceived variability. Mackie et al. (1993) also assessed subjects' response latencies for making variability judgments, and obtained results that were consistent with a retrieval-based process. Although both in-group and out-group judgments appeared to be retrieval based, the in-group judgments took less time than the out-group judgments. This finding contradicts the PDIST model, in that if more in-group than out-group members are retrieved, judgments of in-group variability should take longer.

Given these mixed findings, it is difficult to draw any strong conclusions about the extent to which variability judgments are based on retrieval of exemplars from memory. Moreover, if that process of exemplar retrieval and summation is presumed to be a parallel and implicit process in exemplar models, it may be difficult to obtain evidence for or against these models. That is, one would not necessarily expect to find a relationship between judgment and recall, nor would response latency data necessarily be informative about a process that is assumed not to consume resources and capacity. The issue of whether variability judgments are formed in a strictly retrieval-based fashion will require further research.

C. *Abstraction-Plus-Exemplar Model.* In contrast to the exemplar models just discussed, some researchers have argued that our definition of stereotypes must be broadened to include variability as well as central tendency information. Park and Judd (Judd & Park, 1988; Park & Hastie, 1987; Park et al., 1991) proposed that variability estimates are updated on-line and stored as part of the group stereotype, just like information about any other attribute (Fried & Holyoak, 1984). When required, the stored variability estimate may be retrieved; exemplar retrieval is unnecessary. Particular exemplars also are stored in memory and can be

retrieved to make judgments. However, exemplars are retrieved in addition to, not instead of, pre-stored abstract knowledge. We refer to this model as the abstraction-plus-exemplar model.

According to this model, the out-group homogeneity effect results from differences in the kinds of information used for judging in-group and out-group variability. Whereas judgments of both groups rely on the retrieval of abstract variability estimates, judgments of the in-group (but not the out-group) also are based on some retrieved exemplars. Use of these exemplars should produce higher estimates of variability for in-groups than out-groups.

Park et al. (1991) suggested several reasons why exemplars may play a larger role in in-group than in out-group judgments. First, we simply may be more interested in different kinds of information about in-groups and out-groups (Park & Rothbart, 1982). Specifically, we may be oriented toward defining the typical qualities of out-group members (which focuses on abstract information), but in identifying differences between ourselves and other in-group members (which focuses on in-group exemplar information; Brewer, 1993). Second, the self is more likely to come to mind as an exemplar when making judgments about one's in-group than about an out-group (Park et al., 1991). This may induce comparisons between the self and other in-group members, focusing attention on in-group exemplars both at encoding and retrieval. Third, we may be more motivated to form accurate impressions of in-groups than of out-groups. If so, then we may use a larger sample of exemplars when making variability estimates of an in-group than an out-group. Finally, we may be exposed to different kinds of information about in-groups and out-groups. Our impressions of in-groups are more likely to be based on actual behavioral exemplars that we have witnessed, whereas our knowledge of out-groups often is provided (by media or socializing agents) in the form of general descriptors. We simply may have had less direct experience with individual out-group members (Park et al., 1991; Quattrone, 1986).

Most of the evidence cited as support for on-line processing of variability information comes from failures to obtain relationships between recalled information and variability judgments (Judd & Park, 1988; Park & Hastie, 1987). However, as already discussed, these data are not conclusive and other findings (Mackie et al., 1993) argued against on-line variability abstraction.

However, there is support for some of the other predictions of the abstraction-plus-exemplar model. For example, Park and Judd (1990) obtained positive correlations between self-judgments and in-group judgments, but no relationship between self-judgments and out-group judgments. Also, Park and Judd reported evidence of a relationship between retrieved in-group members and in-group judgments, but no relationship between retrieved out-group members and out-group judgments. These data all support the model's contention that retrieval processes play a larger role in in-group than in out-group variability judgments. On the other hand, when the effects of self and retrieved group members were controlled, the OHE still existed, suggesting that other factors contribute to the OHE (Park

& Judd, 1990). Also, as noted earlier, the Mackie et al. (1993) finding that in-group judgments were made more quickly than out-group judgments questions the assumption of greater in-group than out-group exemplar retrieval.

This model proposes several possible mechanisms that may contribute to judgments of group variability. To date, there is little evidence that variability information is stored as an abstract representation as part of the group stereotype, questioning one of the model's major contentions. Even if abstract information plays little or no role, the model suggests several self-involving and motivational factors that also may influence variability judgments.

d. *Frequency Distribution Model.* Another model has been proposed by Park and Judd (Kraus, Ryan, Judd, Hastie, & Park, 1993; Park, Ryan, & Judd, 1992). The frequency distribution model posits that people spontaneously create mental frequency distributions that summarize the number of group members at different levels of particular attribute dimensions. For instance, subjects may store the number of high-, moderate-, and low-intelligence behaviors performed by group members, or the number of smart, average, or stupid individuals in the group. These frequency distributions themselves do not constitute variability estimates. Rather, when a variability judgment is required, subjects retrieve the distributions and base their variability estimate on the number of levels, or subtypes, used to discriminate among group members along particular dimensions. Thus, a group represented by five different levels (or subtypes) of intelligence will be judged as more variable than a group represented by three levels. In this view, the out-group homogeneity effect occurs because in-group members are more likely to be subtyped than out-group members, due to the same processes differentiating self from other in-group members outlined earlier.

In contrast to the abstraction-plus-exemplar model, this model does not postulate that stereotypes include both variability and central tendency information. The formation of distributions occurs on-line, but variability estimates are based on the retrieval of subtypes. This retrieval process is analogous to PDIST, but rather than focusing on the number of exemplars retrieved, the frequency distribution model focuses on the number of subtypes retrieved. In contrast to PDIST, variability information is attended to on-line, but (in line with PDIST) no abstract variability representation is stored at the group level.

Interestingly, this model portrays a different role of subtyping in stereotype maintenance than typically has been assumed. Traditionally, subtyping has been viewed as protecting stereotypes from change; inconsistent group members are subtyped, leaving the overall group impression intact. However, in this model, the subtyping process may yield greater perceived variability within the group, and therefore undermine stereotyping. This analysis proposes an interesting paradox that is addressed further in the section on stereotype change.

Although only recently introduced, some initial support for the frequency distribution hypothesis has been reported (Kraus et al., 1993; Park et al., 1992).

As predicted, subjects generated more in-group than out-group subtypes, and variability judgments were positively related to the number of subtypes used to classify group members. However, the judgment process hypothesized by the model—that people retrieve subtypes when they make variability judgments—remains untested.

e. *Summary.* Attempts to understand how perceivers understand and make judgments of group variability obviously have produced a growing research literature, and there has been more than a little debate among advocates of the various models we have discussed. In large part, this research has been driven by the continuing puzzle of the out-group homogeneity effect: Why is it that people often perceive greater homogeneity in other groups than in their own? Intuitively the finding is not surprising, but developing a fully satisfactory explanation has not been easy (a task further complicated by recent findings of greater perceived in-group homogeneity; Brewer, 1993; Simon & Brown, 1987; Simon & Hamilton, in press). But the fundamental issue is considerably more than the out-group homogeneity effect; the issue concerns how information about group members is processed, stored in memory, and used in making judgments about the group. That is, it leads us to think about the nature of group representations and what we mean by the term *stereotype*.

As stated earlier, a stereotype pertains to "central tendency beliefs," that is, the attributes thought to characterize a group "as a whole" or "on average." But people certainly retain knowledge of individual group members and of their experiences with them, and perceivers can make at least reasonable estimates of group variability. Moreover, perceived variability is related to a number of effects in stereotyping and intergroup perception (Hastie & Park, 1987; Park et al., 1991) and intergroup behavior (Wilder, 1978). What role, then, do perceptions of variability play in stereotyping?

At present, there is little evidence for on-line abstraction of variability information. However, if variability estimates are generated in a strictly retrieval-based fashion, it becomes unclear what causal role perceptions of variability could play in producing such effects. That is, when performing many of the tasks often related to variability judgments (e.g., generalization tasks), it seems highly unlikely that subjects would spontaneously retrieve information and form a variability estimate prior to responding. If they do not, it is questionable that post hoc, retrieval-based variability estimates would play a causal or mediating role in producing other, related effects. Instead, some other aspect of group information, such as stereotypes or central tendency information, simply could exert direct influence on these responses. That is, it seems possible that variability estimates are merely by-products of stereotyping, rather than contributors to it.

Is such a view tenable? Perhaps variability estimates simply reflect the strength with which a stereotype is held. For example, the more confidence or conviction one has in a stereotype, the more that stereotype might be used in making

judgments of the group, with relatively little reliance on retrieval of exemplar information from memory. In this case, central tendency judgments could be made with relative ease, and the group would be seen as rather homogeneous. On the other hand, for less clearly articulated and established stereotypes, subsequent judgments may rely more heavily on retrieval of exemplars and/or subtypes, which typically might increase perceptions of variability.

Could this view account for the out-group homogeneity effect? Perhaps. Assume that perceivers develop less clearly defined stereotypes of their own group, and hence those self-stereotypes are held with less conviction. In contrast, stereotypes of other groups are formed more readily and are held with greater confidence. Given this difference, judgments of one's own group would draw on relevant exemplar information, whereas judgments of other groups would be based more directly on the stereotype. It would follow, then, that estimates of in-group variability would be greater than would those for out-groups. This view is also compatible with recent research on self-stereotyping, in which both greater self-stereotyping and perceived in-group homogeneity have been found under conditions that would be likely to promote clear and confident self-stereotypes, such as minority status and heightened salience of in-group (Brewer, 1993; Simon & Hamilton, in press).

If variability estimates derive from stereotypes and if other variability-related effects are also products of stereotypes, there certainly would be a positive relationship between these effects and perceived variability. However, to the extent that variability estimates are retrieval based and are made only when they are explicitly requested, we must begin to question the causal or mediational role of perceived variability in producing these effects.

This discussion highlights the importance of understanding the basic representational structure of group-relevant information. Earlier we argued against a pure exemplar-based definition of stereotypes. We questioned why, if categorization and judgment processes can be based exclusively on exemplar information, one would construct a retrieval-based stereotype at all. We summarized evidence that, we argued, supports a conception of stereotypes as abstract representations of, and derived from, group-relevant information. However, if group-relevant information is represented in generalized form, why is variability information not abstracted and stored as part of the stereotype? The answer may be in the relative importance that the perceiver attaches to variability information. That is, the fact that we develop abstract knowledge about some aspects of group-relevant information does not mean that all aspects will necessarily be represented abstractly. Thus, those attributes that are most important for defining a group, or that are most useful for perceivers in making judgments and guiding behavior, are more likely to be abstracted and represented in a stereotype. At this point, it appears that variability information may not be one of those central attributes.

Clearly, there is more work to be done on these issues. The models proposed thus far, and the research they have stimulated, have left many questions unanswered. Nevertheless, they have challenged traditional ways of thinking about the nature of stereotypes, and they have generated new conceptual frameworks that are likely to be the catalyst for additional progress in the years ahead.

[3] STEREOTYPING AND INFORMATION PROCESSING

As mentioned earlier, stereotypes act as expectancies about groups and their members (Hamilton et al., 1990). Like all expectancies, stereotypes guide information processing and often are perpetuated by confirmatory biases that they themselves generate. Indeed, the activation of a stereotype can affect all aspects of social information processing, including attention (Zadny & Gerard, 1974; Bodenhausen, 1988), interpretation (Darley & Gross, 1983; Sagar & Schofield, 1980), inference (Bodenhausen & Wyer, 1985; Krueger & Rothbart, 1988), and retrieval (Bodenhausen & Lichtenstein, 1987; Cohen, 1981; Hamilton & Rose, 1980; Stangor & Duan, 1991). Stereotypes also can influence the types of information perceivers seek about targets (Kunda, 1990; Skov & Sherman, 1986; Snyder & Swann, 1978) and can direct behavior in confirmatory ways, creating self-fulfilling hypotheses (Snyder, Tanke, & Berscheid, 1977; Word, Zanna, & Cooper, 1974).

In this section, we focus primarily on how stereotypes affect encoding, retrieval, and inference processes. We place particular emphasis on topics that have received substantial attention in recent years.

[1] Stereotypes: Encoding and Retrieval Biases

Interpretation Effects. Once activated, stereotypes can affect the interpretation of subsequently presented target information (Darley & Gross, 1983; Duncan, 1976; Kunda & Sherman-Williams, 1992; Sagar & Schofield, 1980). For instance, Sagar and Schofield (1980) found that ambiguous acts were interpreted as more aggressive when performed by a Black target than by a White target. This type of bias is particularly strong when behavioral information is ambiguous. Some researchers have suggested that such interpretational biases occur automatically (Devine, 1989).

In a similar vein, Biernat, Manis, and Nelson (1991) showed that people use different standards of comparison when judging the behaviors of different groups of people. For example, an assertive behavior performed by a woman may be perceived as being more assertive than when that same behavior is performed by a man. Manis et al. argued that the assertiveness of the target person's behavior is evaluated in comparison with the standard of typical assertiveness of the person's gender group, based on stereotypic expectancies. Because people believe that women typically are less assertive, a woman's clearly assertive

behavior is viewed as more assertive than when it is performed by a man, for whom the same behavior is less discrepant from his group standard.

Another interpretation bias involves the attributions perceivers make for stereotype-consistent and stereotype-inconsistent behaviors. Behavior that confirms a stereotype is more likely to be attributed to a target's stable personality factors than disconfirming behavior (Bodenhausen & Wyer, 1985; Macrae & Shepherd, 1989), whereas inconsistent behavior is more likely to be attributed to situational factors. Furthermore, perceivers differentiate between high- and low-credibility sources only when they provide stereotype-inconsistent information (Macrae, Shepherd, & Milne, 1992). There is no preference for high-credibility sources when stereotype-consistent information is provided. Subjects have a more restrictive acceptance threshold for inconsistent information.

Stereotypes also can affect the patterns that are perceived in available information. Hamilton and Rose (1980) showed that subjects overestimated the number of times that stereotypic traits were used to describe members of occupational groups. The groups were described with equal numbers of stereotypic and nonstereotypic terms, so there was no actual relationship between the groups and the terms used to describe them. However, subjects perceived illusory correlations between the groups and their stereotypic traits. In another study in which traits did describe some groups more often than others, subjects were more likely to detect the relationship if the trait-group association was stereotypic than if it was not (Hamilton & Rose, 1980). Thus, stereotypes not only influence interpretation of individual behaviors but also affect the associations that are perceived in patterns of acquired information.

Selective Processing and Recall. Taylor and Crocker (1981) proposed that schemata such as stereotypes function to filter out stereotype-inconsistent information. Accordingly, stereotype-confirming information would receive more attention than inconsistent information (Bodenhausen, 1988; Cohen, 1981; Zadny & Gerard, 1974). Furthermore, this increased attention would cause consistent information to be better incorporated into the perceiver's impression of the target. These processes would yield greater recall of stereotype-consistent information and would serve to perpetuate the stereotype.

Early research suggested a memory advantage for consistent information over stereotype-irrelevant information (Cohen, 1981; Rothbart, Evans, & Fulero, 1979). However, recent analyses of the accumulated research findings showed little evidence of increased encoding or recall of consistent information over inconsistent information (Rojahn & Pettigrew, 1992; Srull & Wyer, 1989; Stangor & McMillan, 1992). On the contrary, most research has shown an advantage for inconsistent information.

However, not all research has been consistent with these general conclusions. Both Stangor and McMillan (1992) and Rojahn and Pettigrew (1992) published extensive meta-analyses of the existing research on memory for expectancy-

congruent and expectancy-incongruent information. Their analyses indicate that, although it is true that there is generally a recall advantage for inconsistent information, this tendency is modified by a number of important variables.

The recall advantage of inconsistent information is due to the increased attention it receives at encoding. Because it violates an expectancy, it is surprising, draws people's attention, and initiates attempts to explain the inconsistency (Clary & Tesser, 1983; Hastie, 1984; Sherman & Hamilton, in press; Stern, Marrs, Millar, & Cole, 1984). In general, the meta-analyses demonstrate that any variable that decreases such inconsistency resolution processes will attenuate recall of inconsistent information, in some cases leading to recall advantages for consistent information.

Studies finding better recall of inconsistent information typically have created trait expectancies in the laboratory, whereas experiments using preexisting stereotypes as expectancies frequently have found better recall for consistent information. There are obviously important differences between trait expectancies and stereotypes. Expectancies pertaining to groups seem to generate less of an inconsistency effect than expectancies pertaining to individuals (Srull, Lichtenstein, & Rothbart, 1985; Stangor & Ruble, 1989; Stern et al., 1984), perhaps because (compared with stereotypes) trait expectancies are stronger and hence yield greater attempts at inconsistency resolution (Srull et al., 1985; but see Stangor & McMillan, 1992, for a different view). In addition, specific traits certainly circumscribe a narrower range of behaviors than stereotypes (Andersen & Klatzky, 1987), leaving less room for inconsistency.

Using preexisting stereotypes, Bodenhausen (1988) showed that subjects had better recall for consistent than inconsistent information. Furthermore, this recall advantage apparently was due to increased attention to consistent information, and not due to interpretational biases or retrieval effects (see also Bodenhausen & Lichtenstein, 1987).

There are other variables that eliminate or even reverse the recall advantage of inconsistent information. Subjects with limited cognitive capacity are less able to engage in inconsistency resolution, and therefore recall consistent information at least as well as inconsistent information (Macrae, Hewstone, & Griffiths, 1993; Srull, 1981; Srull et al., 1985; Stangor & Duan, 1991).

In fact, several researchers suggested that, as capacity saving devices, stereotypes are most likely to bias encoding processes under capacity constraints (Bodenhausen & Lichtenstein, 1987; Macrae et al., 1993; Stangor & Duan, 1991). Under such conditions, the relative ease with which stereotype-consistent information is processed, and the inability to devote appropriate resources to explaining the inconsistent information, will lead to a recall advantage for consistent information. In addition, due to limited capacity, subjects may be unable to form judgments on-line. If so, then not only would the memory load yield increased recall for consistent information, the recalled information also may play an important role in judgments of the target.

Several experiments provide evidence relevant to these points. Using an experimentally induced expectancy, Stangor and Duan (1991) demonstrated that the recall advantage of inconsistent information decreased as the complexity of the processing task increased, presumably depleting capacity (see also Hamilton, Driscoll, & Worth, 1989). In a second study, Stangor and Duan obtained a recall advantage for consistent information under limited capacity conditions, but did not find a relationship between the information recalled and the judgments made about the targets.

Using expectancies based on occupational stereotypes, Macrae et al. (1993) found better memory for inconsistent information under high-capacity conditions and better memory for consistent information under low-capacity conditions. As predicted, they found a relationship between recall and judgment only under limited capacity conditions. Similarly, Bodenhausen and Lichtenstein (1987) showed a recall advantage for consistent information under anticipated high load conditions. However, like Stangor and Duan (1991), they also found no relationship between recall and judgments.

Summarizing to this point, inconsistent information generally is recalled better than consistent information, but certain conditions favor recall of consistent information. This latter pattern of results appears to be more likely when the expectancy pertains to a group (e.g., a stereotype) than when it pertains to an individual, and when capacity is limited (although it appears that recalled information may not affect judgments). Most researchers attribute this pattern to subjects' decreased ability to process inconsistent information (Macrae et al., 1993; Srull, 1981; Srull et al., 1985; Stangor & Duan, 1991). Without increased attention during encoding, inconsistent information is recalled less well than consistent information.

To this point, only Bodenhausen (1988) has demonstrated that consistency biases may involve more than equalizing the attention paid to consistent and inconsistent information. He provided evidence that consistent information actually receives more attention than inconsistent information. Given that there is frequently a lack of a relationship between judgments and recall (e.g., Bodenhausen & Lichtenstein, 1987; Stangor & Duan, 1991), Bodenhausen's findings may be particularly important. Stereotype-driven biases that favor the recall of consistent information may seem to be of little consequence if that information does not affect subsequent judgments. However, if such biases affect the attention given to consistent information, they will affect both on-line judgments (by affecting the weight given to different information) and recall of information (well-attended information tends to be recalled better). Thus, establishing such stereotype-congruent attentional biases has important implications.

In addition, there are other reasons that inconsistent information may have little influence on judgments. Inconsistent information often is explained away or attributed to situational causes as it is encoded. If so, then although this incongruency-resolution process would lead to a recall advantage for inconsistent information, it also would decrease the impact of that information on subsequent judgments. On the other hand, if subjects are unable to process inconsistent information in this way at encoding (e.g., under a high memory load), they are less likely to form on-line judgments and inconsistent information is unlikely to be well retrieved. Consequently, subsequent judgments would be memory-based, but would not be influenced heavily by the recall of inconsistent information. In both of these cases, then, the inconsistent information would not have strong impact on judgments.

Response Biases. In their meta-analyses, both Stangor and McMillan (1992) and Rojahn and Pettigrew (1992) found a strong response bias for consistent information. That is, recognition measures that are not corrected for guessing demonstrate a strong advantage for consistent information. This suggests that subjects frequently guess that they have seen expectancy-consistent information when they have not. Such strategies may override any effects due to increased encoding and recall of inconsistent information.

In a similar vein, Slusher and Anderson (1987) demonstrated that, with the passage of time, it becomes difficult for subjects to distinguish between stereotype-consistent information that actually was presented and stereotype-consistent inferences that subjects drew from that information. Subjects sometimes are unable to distinguish what they know from what they believe.

Summary. Stereotypes can serve to bias information processing in a number of self-perpetuating ways. Stereotypes can bias the initial interpretation of information in a confirmatory fashion. Attributions made about consistent and inconsistent behaviors also serve to confirm the stereotype. Illusory correlations between stereotypical information and targets inflate the perceived prevalence of consistent information. Under appropriate conditions, stereotype-consistent information may receive more attention than and be better recalled than inconsistent information. Furthermore, stereotype-consistent information may be recognized falsely as having been encountered, yielding a general response bias for consistent information. In total, these various processes demonstrate that one of the primary functions of stereotypes is self-maintenance.

Stereotyping and Inference

One important effect of group stereotypes is that they color our perceptions of individual group members. The stereotype allows us to infer an individual's characteristics based on group membership. Therefore, stereotypes should reduce the impact of group members' personal behaviors and attributes in our impressions of them (Bodenhausen & Wyer, 1985; Krueger & Rothbart, 1988).

However, this is not invariably the case. Some circumstances favor the use of individuating information over stereotypes. In their now-famous studies, Locksley,

Borgida, Brekke, and Hepburn (1980) demonstrated that judgments about the assertiveness of male and female target persons were based on the nature of the behaviors they performed, and not on their gender. This was true even when subjects learned only a single diagnostic behavior about the target. Only when no other diagnostic information was available did gender stereotypes affect subjects' ratings of the targets. This led Locksley et al. to suggest that stereotypes only will be utilized when the perceiver has no individuating information about the target person.

More recent research indicates that this conclusion was somewhat overstated. In a series of studies, Krueger and Rothbart (1988) demonstrated that individuation depends on both the diagnosticity of the stereotype and the diagnosticity of the individuating information. In their first experiment, subjects read only a single individuating behavior. In contrast to the Locksley et al. results, even when the behavior was highly diagnostic, stereotypes exerted strong influence on judgments. However, in subsequent experiments, stable behavioral and trait information did override stereotype effects. The key to overriding the stereotypes was the presence of temporal consistency in targets' behaviors. Such consistency is far more diagnostic than a single behavior. Other researchers have documented that individuation occurs only when highly diagnostic information about individuals is provided (Heilman, 1984; Kunda & Sherman-Williams, 1992; Krueger & Rothbart, 1988). How, then, do we explain the findings from the Locksley et al. (1980) studies? Krueger and Rothbart (1988) argued that the circumstances in these studies were just right for individuation to occur, in that the behavioral information was highly diagnostic and gender stereotypes may be weak predictors of assertive behavior (the target behavior used by Locksley et al.).

An experiment by Darley and Gross (1983) demonstrates what may happen when stereotypes are activated in the presence of ambiguous information. Interestingly, whereas Locksley et al. (1980) found that stereotyping occurred only in the absence of specific information, in Darley and Gross's study, stereotyping occurred only when subjects received specific target information. Subjects who learned about a target person's socioeconomic status (SES) and received no behavioral information did not use their SES stereotype in evaluating the person's academic ability. However, subjects who received behavioral information in addition to SES information did differentiate high and low SES students. Analyses revealed that these subjects processed the behavioral information in a hypothesis-confirming manner. Specifically, behavior that was consistent with subjects' expectancies was recalled better and weighted more heavily, and ambiguous behaviors were interpreted differently depending on the target's SES.

Similarly, there was some evidence from Locksley et al. (1980) that subjects interpreted behavior in a biased manner. Assertive behaviors were rated as more masculine than feminine, and passive behaviors were rated as more feminine than masculine. Whereas male and female targets may have been rated as equally assertive, it is not clear that subjects perceived the targets' behavior to be equally masculine or feminine. As noted earlier, Biernat et al. (1991) reinterpreted these results by proposing that people use different standards of comparison when judging males and females.

The findings of Locksley et al. (1980) and Krueger and Rothbart (1988) suggest that stereotypes can be overcome if individuating information is strongly diagnostic. Similarly, other models suggest that if individuating information clearly contradicts a stereotype, then individuation may occur (Brewer, 1988; Fiske & Neuberg, 1990).

Most current theorists predict that stereotypes have clear priority over individuating information in people's impressions of others (Bodenhausen & Wyer, 1985; Brewer, 1988; Devine, 1989; Dovidio, Evans, & Tyler, 1986; Fiske & Neuberg, 1990). This prediction is based, in part, on the simple fact that category-based stereotypes are more salient and easier to utilize than individuating information. Because of capacity requirements, these models suggest that individuation will only occur if perceivers are motivated to attend to target persons carefully. For instance, if perceivers are motivated by personal relevance or desire accuracy, they may individuate target persons (Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990). Also, evidence shows that forming an individuated impression (Fiske, Neuberg, Beattie, & Milberg, 1987) takes longer and requires greater attentional resources (Bechtold et al., 1986) than does a stereotype-based impression.

Stereotypes as Capacity-Conserving Devices. If it is true that stereotypes function to preserve cognitive resources, then they should be particularly useful in conditions of limited capacity. Recently, a number of studies examined stereotype use under high- and low-capacity conditions. In one study (Bodenhausen & Lichtenstein, 1987), when subjects performed a cognitively demanding judgment task, they produced stereotypical judgments and a recall advantage for stereotype-consistent information. This was not the case when they performed a relatively simple judgment task.

Macrae et al. (in press) provided stronger support for a capacity-saving function for stereotypes. Subjects were asked to form impressions of target individuals while performing an audiotape-monitoring distractor task. To examine the capacity-preserving function of stereotypes, some subjects were given stereotype labels (e.g., skinhead) along with the information about the target. Presumably, this information would simplify the impression formation task by providing subjects with a stereotypic theme to guide their impressions.

The results showed that subjects who were given the stereotype information recalled more of the traits used to describe the target than did subjects who were not given the stereotype. However, this was only true for stereotype-consistent traits. Thus, under capacity constraints, stereotypes function as organizational structures, facilitating the learning of stereotype-consistent information. These subjects also performed better on a test of the material presented in the tape-

recorded distractor task. Again, this suggests that the stereotypes functioned to preserve capacity required for the impression task, such that more resources could be directed toward the tape-monitoring task. Macrae et al. (in press) also reported additional experiments confirming this interpretation. Gilbert and Hixon (1991) also showed greater reliance on stereotypes when making judgments under high cognitive load.

Additional support for viewing stereotypes as capacity-preserving devices comes from some ingenious studies of the use of stereotypes depending on subjects' circadian arousal levels (Bodenhausen, 1990). People go through daily variations in their arousal levels that affect capacity and efficiency of working memory. Some people reach their functional peak during the morning ("morning people"), and some in the evening ("night people"). Bodenhausen showed that subjects relied more on stereotypes when they were at the low end of their circadian cycles. Morning people used stereotypes more in the evening and night people used them more in the morning. This finding suggests that stereotypes are utilized as capacity-saving devices when people are incapable (or unwilling) to process optimally.

Together, the results of these experiments provide solid evidence that stereotypes function to preserve cognitive resources.

(3) The Automaticity of Stereotyping

Some researchers have argued that stereotyping is not only easier than individuation, but that it is automatic as well (Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990). Why might stereotypes be activated automatically upon categorization? One reason is that many stereotypes apparently are learned at a very early age (Katz, 1976). As a consequence, stereotypes have a long history of activation and are likely to be highly accessible (Higgins & King, 1981). In addition, due to their long-term usage, stereotypes are likely to become somewhat proceduralized (Smith, 1990).

Several recent studies have examined the automaticity of stereotyping activation and application. For example, Devine (1989) demonstrated that nonconscious priming of category labels (e.g., Blacks, Negroes) and terms related to the stereotype of Blacks (*poor*, *slavery*, *athletic*) affected subsequent ratings of a race-unspecified target. Devine argued that these results were due to the activation of the Black stereotype in response to the primes. However, it is somewhat unclear whether her results were due to stereotype activation or to simple semantic priming (Bargh & Pietromonaco, 1982).

Macrae et al. (in press) modified their study on the capacity-preserving functions of stereotypes (described earlier) to test whether these effects occurred automatically. Subjects performed both an impression task and a tape monitoring task at the same time. Again, some subjects received stereotype labels to assist them in the impression formation task. However, in this study, the stereotype labels

were presented subliminally, outside of subjects' conscious awareness. Their results showed that subjects who were exposed subconsciously to the stereotype labels both recalled more of the traits and performed better on the test regarding the recorded material. These results support the idea that the stereotypes allowed these subjects to conserve resources on the impression task that could then be devoted to the tape monitoring task. They also show that this capacity-preserving function of stereotypes is not dependent on conscious access to the stereotypes. The stereotypes apparently were activated and used automatically, without conscious awareness.

Banaji and Greenwald (in press) presented further evidence that stereotypes can influence perceptions and judgments without intention or awareness. They modified an experiment performed by Jacoby, Kelley, Brown, and Jasechko (1989) demonstrating unconscious influences on memory. Jacoby et al. had found that subjects judged names that they previously had seen (but could not remember seeing) as famous. Those names apparently were somewhat familiar to subjects, even though they could not explicitly remember them, leading to their becoming "famous overnight." Banaji and Greenwald modified this procedure by making half of the names male and half of the names female. Subjects' tendency for judging previously seen (but unremembered) names to be famous was significantly greater when those names were male than when they were female, suggesting an implicit stereotype that associates males more closely with fame than females.

Research has shown that the presentation of trait information in one context can influence subsequent judgments of an ambiguously described target in an unrelated context (Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979). Banaji and Greenwald (in press) also showed that the effect of a prime depended on its stereotypicality for the judged target. For example, when the trait "dependent" (stereotypic for females, counterstereotypic for males) was used as a prime, subjects subsequently judged a female target person as more dependent and a male target person as less dependent, compared with subjects exposed to a neutral prime. Thus, the stereotypicality of the match between the prime and the target can affect the extent and nature of priming effects. These are implicit, unconscious effects because they occur without subjects being aware of the influence of the prime. Together, the results of these studies suggest that gender stereotypes can affect information processing and judgments subconsciously.

Contrary to these research findings, Gilbert and Hixon (1991) reported results suggesting that stereotypes are not activated automatically. Subjects under a low cognitive load, when exposed to an Asian target person, completed word fragments in a stereotype-biased fashion, suggesting that seeing the target person had activated the stereotype. However, subjects under a high cognitive load did not produce stereotypic word-fragment completions. Because the cognitive load manipulation affected performance in this way, stereotype activation apparently was not automatic, but was subject to capacity limitations.

In this study, subjects under high cognitive load in the activation stage were

able to identify the target person's race correctly. Therefore, although stereotype activation may not have occurred in this condition, it was not because subjects failed to categorize the target person. This is a potentially important point. Some models assume that categorization is an automatic process based on perceptible features (Brewer, 1988; Devine, 1989; Fiske & Neuberg, 1990) and that, upon categorization, stereotypes are activated automatically. However, Gilbert and Hixon's results suggest that categorization can occur without stereotype activation.

Summary. Although the findings are not entirely consistent, evidence is accumulating that stereotypes are activated and applied routinely under some conditions. Stereotypes appear to serve a capacity-preserving function. They are particularly likely to be used under low capacity (Bodenhausen, 1990; Gilbert & Hixon, 1992; Macrae et al., in press) or low motivation (Neuberg & Fiske, 1987). On the other hand, if subjects have sufficient capacity and motivation, and if individuating information is particularly diagnostic or disconfirmatory, stereotypes may be overridden.

Just as stereotyping requires less effort than individuation, categorization may require less effort than stereotyping. Categorization appears to be a necessary, but not a sufficient, precondition for stereotype activation and use (Gilbert & Hixon, 1992).

[4] AFFECT, COGNITION, AND STEREOTYPING

If there is one domain of social perception where we might expect affective processes to play a particularly important role, it would be intergroup perception. Historically, emotion and feelings were viewed as major elements contributing to stereotype-based biases and to the tenacity with which those stereotypes are held (Allport, 1954; Katz, 1981; Lippmann, 1922). Yet, for over a decade, beginning in the mid-1970s, stereotype research was almost totally focused either on the cognitive underpinnings and functioning of stereotypes or on the role of social structures and mass media in the development and maintenance of these belief systems. Discussions of the stereotype literature during this period (Hamilton, 1979; Hamilton & Trolier, 1986; Hewstone, 1989; Stephan, 1985; Stroebe & Insko, 1989) included relatively little coverage of the role of affective factors in stereotyping and intergroup perceptions. However, there were occasional calls for greater attention to affective processes and their interrelations with cognitive processes in this domain (Fiske, 1982; Hamilton, 1981).

Fortunately, in recent years, we have witnessed not only a resurgence of interest in affect, emotion, motivation, and related processes per se, but also (and more importantly) a focus on the interactive influence of affect and cognition on social information processing and judgment processes. This development has impacted research on intergroup perceptions and has illuminated a variety of ways

in which affect and cognition interface in intergroup perceptions. In this section, we highlight some of these recent developments. A broader and more thorough coverage can be found in a recent edited volume focused specifically on the interactive roles of affect and cognition in intergroup perception (Mackie & Hamilton, 1993).

In conceptualizing how affect might influence intergroup perception, it is useful to distinguish between different sources of affect in the intergroup situation. The affect experienced by the perceiver might have been generated prior to and independently of the intergroup context. Bodenhausen (1993) referred to this as incidental affect, and in this case we typically are concerned with the influence of the perceiver's arousal or mood state on the formation and/or use of stereotypes in processing information about group members. Alternatively, the affect experienced by the perceiver may have its origins in the intergroup context, being generated by one's perceptions of or interaction with members of some outgroup. In this case, the affect is integral to the intergroup context (Bodenhausen, 1993). Here the perceiver's preexisting stereotype, as well as other aspects of the intergroup situation, can be the catalyst for the experienced affect. In addition, the affect generated in this setting can influence subsequent information processing, which in turn can affect further the participant's feelings, perceptions, and behavior. In the following sections, we discuss recent research on the effects of both incidental and integral affect in intergroup contexts.

(1) Incidental Affect and Stereotyping

One way that affect influences stereotyping is through the effect of the perceiver's mood on the way information about groups and group members is processed. Research has shown that mood states can have both cognitive and motivational effects on social information processing (Forgas, 1992; Isen, 1984; Mackie, Asuncion, & Rosselli, 1992; Mackie & Worth, 1991; Schwarz, 1990). However, in any given situation, the motivational and cognitive consequences of mood might suggest different effects on information processing. From a motivational perspective, positive and negative moods have differing implications for information processing. Positive mood is predicted to decrease the thoroughness with which information is processed, either to prevent that information from distracting one away from a pleasant affective state or because the positive mood conveys the message that, because "all is well," there is no need to focus on new material. Negative mood would stimulate more extensive processing of new information, either as a means of distracting oneself from the negative affect or because the negative mood informs the person that the current situation needs to be altered. From a cognitive perspective, mood influences processing by interfering with one's ability to perform cognitive tasks. This interference may emanate from the mood state distracting one's attention from the task at hand, or it may result from reduced cognitive capacity due to the mood consuming some cognitive resources.

In either case, positive and negative moods should have much the same effect on these processes.

Research by a number of investigators has demonstrated motivational and cognitive effects of mood on both the formation and use of stereotypes. In these studies, subjects typically are induced into a positive, neutral, or negative mood state (e.g., by watching affectively toned videotaped material or by thinking about pleasant or unpleasant past experiences), and then, in a presumably separate experiment, they are asked to process information pertinent to certain social groups. The effect of the induced mood on subjects' judgments, recall, processing times, and other relevant measures can then be determined.

Some studies have examined the effects of mood on processes that would contribute to the initial development of stereotypic responding. For example, Stroessner and Mackie (1992, 1993) showed that subjects in positive and negative (compared with neutral) mood conditions were less able to detect the degree of variation among members of a group, and hence underestimated overall group variability. As noted in an earlier section, insensitivity to within-group differences enhances the likelihood that perceivers will generalize across group members in attributing characteristics to the group (i.e., stereotyping). In the same vein, Wilder (1993; Wilder & Shapiro, 1989) showed that anxious subjects were less likely to differentiate among members of a group and assimilated a deviant member to the group as a whole. Similarly, research using an illusory correlation paradigm (Stroessner et al. 1992; Hamilton, Stroessner, & Mackie, 1993) showed that both positive and negative mood states (compared with a neutral mood condition) reduced subjects' sensitivity to the distinctiveness properties of the stimulus information (i.e., the relative infrequency of the minority group performing undesirable behaviors), and hence these subjects did not form an illusory correlation. In all of these cases, induced mood undermined the thoroughness with which group-descriptive information was processed. Interestingly, in one case, this effect created conditions increasing the likelihood of stereotyping (i.e., reduced detection of within-group variability), whereas in the other case, the effect decreased the likelihood of stereotype formation (i.e., through undermining the illusory correlation effect).

Other studies investigated the effects of incidental affect on the use of existing stereotypes in perceptions of social groups (Bodenhausen, 1993; Forgas & Moylan, 1991; Hamilton et al., 1993; Kim & Baron, 1988; Mackie et al., 1989). One illustration is found in some recent work by Bodenhausen (1993). As discussed earlier, Bodenhausen (1988, 1990; Bodenhausen & Lichtenstein, 1987; Bodenhausen & Wyer, 1985) has argued that stereotypes can be used as a heuristic or simplified processing strategy under conditions of limited processing capacity or reduced motivation. Drawing on research showing that affect has those effects on cognitive processing, he reasoned that induced mood might lead to increased reliance on stereotypes in social judgments.

In one study testing this hypothesis, Bodenhausen (1993) first induced happiness, sadness, or anger and then had subjects read a description of a case in which one student physically assaulted another. For half of the subjects, the accused student's name identified him as Hispanic, in which case the aggressive act was consistent with the stereotype of his ethnic group. The subjects' task was to evaluate the student's guilt. As predicted, subjects in happy and angry moods—affects hypothesized to reduce processing capacity and motivation—judged the case in stereotypic terms. That is, subjects in these conditions who read about a stereotyped defendant were more likely to find him guilty than were subjects reading about a student without a Hispanic surname. In contrast, this difference did not occur for subjects in a sad mood, consistent with Bodenhausen's (1993) prediction that sadness does not constrain processing and hence would not lead to reliance on a group stereotype. (Although these results followed from Bodenhausen's analysis, the issue of under what conditions sadness leads to more or less careful processing is a matter of debate; see Mackie & Worth, 1991; Stroessner et al., 1992; Stroessner & Mackie, 1992).

As these studies illustrate, incidental affect can have consequences for both the nature and extent of the perceiver's processing of group-relevant information. The ramifications of these effects for both stereotype formation and the use of existing stereotypes in social judgments will continue to be the focus of research activity.

(2) Integral Affect and Stereotyping

Incidental affect refers to prior feeling states that influence the perceiver's perceptions, because he or she brings them to the intergroup context. Therefore, incidental affect always precedes its effect on stereotyping, and the question concerns how the affect influences cognitive processing in this context. The interplay between cognitive and affective factors involved in stereotyping becomes even more complex when the affect is generated in and is integral to the intergroup context. In this situation, cognitive and affective factors each can have a causal effect on the other. That is, cognitive factors can initiate affective reactions, and affective reactions in turn can influence the nature of cognitive processing. We illustrate these points with some recent research examples.

As discussed earlier, participants in many intergroup contexts react in terms of a distinction between in-group and out-group; we already discussed research relevant to several topics (e.g., in-group bias, out-group homogeneity effects) deriving from that distinction. Recently, Dovidio and Gaertner (1993) argued that in-group/out-group categorization automatically activates category-based affective reactions that directly influence perceptions of and expectations about interactions with group members. For example, they have shown that subliminal priming with terms like *we* and *they* automatically generates positive and negative reactions that have evaluative consequences for otherwise neutral stimuli. These reactions also include differing expectations for how pleasant their interactions with target persons would be.

Affective consequences of intergroup perception extend beyond mere in-group/out-group categorization. Whereas the stereotype literature typically has focused on the trait-like beliefs activated by perception of group memberships, several recent studies have shown that the stereotypes associated with various social groups can generate affective reactions as well (Dijker, 1987; Jackson & Sullivan, 1989; Vanman & Miller, 1993). For example, Dijker (1987) showed that his Dutch subjects had distinct patterns of emotional reactions (as measured by rating scales) to three ethnic out-groups living in the Netherlands. Moreover, Vanman and Miller (1993) demonstrated that psychophysiological measures could detect different patterns of facial muscular activity, reflecting positive and negative affective reactions, as they viewed members of different groups. Several studies found that these affective reactions are as (and sometimes more) effective predictors of intergroup attitudes as are stereotypic beliefs (Esses, Haddock, & Zanna, 1993; Stangor, Sullivan, & Ford, 1991; Stephan & Stephan, 1993).

Other research has shown that intergroup interaction (real or anticipated) can generate affective responses that can influence the course of social interaction. Stephan and Stephan (1985) reviewed evidence supporting their hypothesis that intergroup contexts often generate anxiety that can lead to increased stereotyping, disrupted interaction, and avoidance of future interaction. Greater experience in intergroup contact may or may not alleviate these tendencies. Stephan and Stephan (1985) found that increased intergroup contact was correlated negatively with these effects of intergroup anxiety, but Dijker (1987) found stronger (negative) emotional reactions to out-groups among those who had had more interaction with them. Obviously, other factors can moderate these relationships. One relevant factor is the nature of the interaction context. For example, Wilder (1993; Wilder & Shapiro, 1989) found that competitive interaction generated anxiety-related effects on intergroup perceptions, whereas noncompetitive interaction did not. Similarly, Fiske and Ruscher (1993) argued that the nature of the interdependence relationship between group members can influence the extent to which out-group members are seen as disrupting the perceiver's goal attainment in the interaction (see also Vanman & Miller, 1993).

Can these effects of integral affect be overcome? Earlier, we discussed Devine's (1989) hypothesis that cultural stereotypes of prominent ethnic groups are activated automatically when group members are encountered, and that low-prejudiced persons must counteract those reactions by applying their personal (nonprejudiced) beliefs about those groups. More recently, Devine and her colleagues (Devine & Monteith, 1993; Devine, Monteith, Zuwerink, & Elliot, 1991) have investigated the affective consequences of the discrepancy between people's beliefs about how they would and how they should respond in intergroup contexts. Specifically, they argued that the low-prejudiced person's failure to inhibit stereotypic responses to an out-group member activates a discrepancy between internalized nonprejudiced standards and actual responses. This discrepancy, which threatens the person's nonprejudiced values and self-concept, produces negative

affective responses (e.g., guilt, self-criticism), which in turn can lead to efforts to control stereotypic responding. Through this process, persons can learn to respond in ways consistent with their nonprejudiced personal beliefs and standards.

Most research on stereotype use has focused on processes that transpire in the stereotype holder and examine the cognitive and affective consequences of those beliefs. Much less research has considered the cognitive and affective dynamics of the target of stereotypic beliefs in social interaction. Recently, however, Crocker and Major (1989; Major & Crocker, 1993) have investigated the affective consequences of stereotypes for the way such target persons experience social interactions. Specifically, they showed that stereotyped target persons experience attributional ambiguity in interpreting the causes of others' behavior toward them. For example, a negative evaluation of a minority person's performance received from a majority group member might be due to a valid assessment of the target person's poor performance. Alternatively, it could be attributed to the prejudicial attitudes of the evaluator. In the latter case, attributing the negative feedback to the evaluator's prejudice buffers the affective impact of that feedback, and thereby protects one's self-esteem. Similarly, a favorable evaluation might constitute legitimate praise for commendable performance or, alternatively, it might reflect the evaluator's effort to avoid appearing prejudiced. In this case, the positive effects of favorable evaluation would be muted, again diminishing the effect of feedback on one's affective reactions and self-esteem. As a consequence, both positive and negative feedback may have diminished affective impact on minority group members. These authors reported a series of experiments illuminating the conditions that both augment and diminish this attributional ambiguity and its effects on self-evaluations.

The research reviewed here illustrates the variety of ways that affective and cognitive factors intermingle in determining social perceptions, emotional reactions, and behavioral dynamics in intergroup contexts. This work effectively highlights the important and complex ways that stereotypes can influence, and be influenced by, affective concomitants in intergroup settings.

[5] STEREOTYPE CHANGE

How can we change people's stereotypes? Given the pervasiveness of stereotypes and the frequency with which they unfairly victimize members of ethnic groups, it becomes especially important to identify means by which stereotypic beliefs can be undermined or modified. Unfortunately, despite the importance of this question, the problem of changing stereotypes remains a very real and unsolved dilemma.

Although the three conceptual orientations to stereotyping focus on different underlying mechanisms, they all share one common element relating to this issue. In all three cases, the mechanisms that are emphasized to explain why stereo-

types exist and persist are the reasons why stereotypes are difficult to change. We briefly illustrate this point for all three perspectives.

In the psychodynamic approach, stereotypic beliefs about out-group members serve intraphysic needs, and hence are functionally important to the inner workings of the perceiver's mentality. For example, in classical psychodynamic theory, stereotyping reflects the operation of defense mechanisms such as projection and displacement, mechanisms that are functional in keeping unacceptable sources of anxiety out of conscious awareness. Changing stereotypic beliefs might undermine the effectiveness of these defensive processes, thereby creating new sources of anxiety. Similarly, for more contemporary motivational theories that emphasize the self-esteem benefits of viewing out-groups in negative terms, changing stereotypic beliefs about out-groups would remove a potentially important mechanism contributing to the person's positive self-regard. In both cases, stereotyping serves an important function within the person's overall personality. Consequently, simply presenting the person with evidence that some belief is wrong would hardly be effective. Thus, for these approaches, attempts to change stereotypes somehow must confront the question of how these needs are to be met and how these functions are to be served by alternate means. Otherwise, strategies toward change are likely to produce effects that are transitory at best.

Similar difficulties confront efforts toward change from the sociocultural perspective. This viewpoint emphasizes the importance of the social environment (authority figures, peer groups, social roles, media portrayals, etc.) in the formation and maintenance of stereotypic belief systems. These beliefs are acquired through social learning and are maintained through social reinforcement. Again, simply presenting the person with evidence that some belief is wrong hardly would be effective. Attempts to change the person's beliefs somehow would require altering elements of the person's social environment, if not the social structure in which the person lives. Otherwise, the belief-maintaining supports would still be in place.

For the cognitive approach, the story is much the same, although for different reasons. The importance of the categorization process in social perception seems fundamental, both for identifying the regularities and differences in one's social environment and for simplifying the task of information processing in a complex stimulus world. Once a pattern of beliefs becomes associated with those categories, the resulting stereotypes often guide subsequent information processing. We have discussed numerous examples of how existing stereotypes can influence the way new information about groups and group members is processed. In virtually every case, the resulting bias serves to maintain the status quo. The perceiver is more likely to attend to and notice stereotype-consistent information, to make stereotype-consistent inferences, to recall stereotype-consistent information, and so on. The overall consequence is that the perceiver "sees" a pattern of information that seems to provide evidence for the "validity" of the beliefs that themselves influenced the way the information was processed. Once again, simply presenting evidence that a particular belief is wrong seems unlikely to induce effective change.

Thus, all three conceptual orientations effectively point to mechanisms that serve to maintain preexisting stereotypic beliefs, making change difficult. Yet, despite their persistence, we know that stereotypes do change, at least under certain conditions.

In this regard, it is interesting that the same research that often is cited as documenting the persistence of stereotypes over time also presents evidence of dramatic change in stereotypes. This evidence comes from the well-known "Princeton trilogy," in which the same stereotype assessment procedures were used in three studies spanning some 35 years—Katz and Braly (1933) during the depression; Gilbert (1951) in the post-World War II years; and Karlins, Coffman, and Walters (1969) in the 1960s. In these studies, subjects were asked to indicate, on an adjective checklist, those attributes that were most characteristic of each of several national and ethnic groups. Attributes most commonly endorsed were taken to constitute the primary features of the stereotype of that group. One prominent aspect of the findings of these studies is the remarkable stability of some of these stereotypes over a period of three and a half decades. For example, at each time period, the most frequently endorsed characteristics of Negroes included lazy, happy-go-lucky, and musical. However, this stability is only part of the story from this research. The most prominent features of the stereotype of Japanese changed from intelligent, industrious, and progressive in 1933 to imitative, sly, and extremely nationalistic in 1951, and then to ambitious, efficient, and loyal to family ties in 1967. The reason for these substantial shifts can be attributed, of course, to the intervention of World War II between the first and second assessments, and to Japan's subsequent reconstruction and realignment with the United States between the second and third assessments. Similar, although perhaps less dramatic, changes have occurred in stereotypes of African-Americans and women as a function of the civil rights and women's movements, both of which gained momentum during the 1960s. These findings document that stereotypes do change under some conditions. The problem then becomes to identify those mechanisms that are important in causing and in mediating these changes.

Models of Stereotype Change

Although our understanding of this issue is still far from adequate, social cognition research in the last decade has at least provided some conceptual bases for analyzing this process. These analyses focus on how information that disconfirms a stereotype might affect the preexisting stereotypic beliefs and when that effect is most likely to occur. Several models or mechanisms by which such change might take place have been proposed. The bookkeeping model (Rothbart, 1981) suggests a gradual course of change as increasing amounts of disconfirming information are encountered and accumulated. According to this view, each new

piece of information—confirming or disconfirming—acquired about a target group leads to an adjustment in the stereotypic beliefs, either strengthening or weakening the existing stereotype by some modest amount. Hence, stereotypes are constantly in a state of on-line revision as new information is incorporated. In this process, change would occur slowly as current beliefs are gradually adjusted in light of recent information and experience.

A second process, which Rothbart (1981) called the conversion model, posits that change can occur suddenly and significantly in response to critically important and convincing (to the perceiver) instances of stereotype disconfirmation. That is, some experience with out-group members, some challenging new piece of information about the group, or some group-relevant political or cultural event would be sufficiently compelling to lead the stereotype holder to substantially revise his or her beliefs about the target group. In contrast to the bookkeeping model's slow, gradual change in response to accumulating bits of information, this process involves a rapid change of major magnitude (i.e., a conversion) in response to a dramatic disconfirmation experience.

A third model of change, derived from cognitive research on category representations, focuses on a subtyping process (Brewer et al., 1981; Taylor, 1981). In this view, disconfirming information about members of a stereotyped group results in the differentiation of the large group (superordinate category) into subtypes (subordinate categories), with separate patterns of beliefs associated with each of the subtypes. In essence, the process is one of breaking down the large group into smaller subcategories. In contrast to the previous models, this process does not necessarily involve any actual change in preexisting beliefs, but instead posits a mechanism by which those initial beliefs would be applied less generally.

(2) Research Evidence

a. *Patterns of Disconfirming Information.* The three models just described discuss alternate routes by which disconfirming information might bring about change in stereotypic beliefs. The situation addressed by all of them reflects important elements of everyday experiences. That is, we begin with a person who holds a stereotype about a particular target group. Then, over time and through various experiences, the person encounters a number of individual members of that target group, some of whom behave in a manner consistent with stereotype-based expectancies, whereas the behavior of others seems to violate those expectancies. In essence, the perceiver faces an information processing task in which the available information provides some mixture of confirmation and disconfirmation of prior beliefs. Given this situation, what properties of that "mixture" are more or less likely to induce stereotype change, and why? This question has been the focus of some recent research.

Several studies have used essentially the same paradigm to study this issue: Subjects learn information about several members of a stereotyped group, each

person being described by several items of information. Each information item either confirms or disconfirms the stereotype, or is unrelated to the stereotype. The manipulation of primary interest is whether the disconfirming items are (a) concentrated in describing one (or a small proportion) of the group members or (b) dispersed throughout the descriptions of all (or most) group members. The question of interest concerns how these patterns of disconfirming information affect ratings of the group on stereotypic attributes.

The results of these studies are not entirely consistent, and they have been interpreted by various authors as providing evidence in support of all three of the models discussed previously. For example, some studies (Gurwitz & Dodge, 1977; Hewstone, Johnston, & Aird, 1992) found less stereotypic ratings when the disconfirming information was concentrated in the descriptions of one or a few individuals. Because these individuals provide pronounced or dramatic violation of stereotypic expectancies, this evidence has been interpreted as support for the conversion model of stereotype change. On the other hand, other studies (Johnston & Hewstone, 1992; Weber & Crocker, 1983) obtained less evidence of stereotyping when the disconfirming items were dispersed across several group members than when concentrated in a few. The greater effectiveness of disconfirming information under dispersed than concentrated conditions has been interpreted as support for the subtyping model. The argument here is that, when disconfirming items are concentrated in a few group members, those individuals will be grouped together in a subtype, and therefore their disconfirmation value will have relatively little effect on perceptions of the group as a whole. In fact, hierarchical clustering analyses indicate that disconfirmers are subgrouped more clearly when the disconfirming information is concentrated in a few members (Johnston & Hewstone, 1992).

Research also indicates that these general patterns may be dependent on other factors, such as the number of group members presented to the subjects (Weber & Crocker, 1983, Experiments 1 & 4), the number of groups presented (Weber & Crocker, 1983, Experiment 4), and whether the stereotypic expectancy about the target group implies that group members are homogeneous or heterogeneous (Hewstone et al., 1992). It also remains unclear whether these patterns of disconfirming information have their primary effect on perceptions of stereotype-consistent traits (Hewstone et al., 1992), stereotype-inconsistent traits (Johnston & Hewstone, 1992), or both (Weber & Crocker, 1983).

b. *Subtyping and Stereotype Change.* The finding that stereotype-disconfirming information, particularly when concentrated in a few group members, can result in the creation of subtypes that are characterized by distinct sets of attributes demonstrates that people develop hierarchical systems of social categories. As these subtypes develop, stereotypes may form pertaining to them as well as to the superordinate category (Brewer et al., 1981; Devine & Baker, 1991). Given that disconfirming members can form a subtype, the question then arises: What

is the relationship between subtyping and stereotype change? That is, how does the formation of subtypes affect the stereotype of the overall group? Subtyping commonly is cited as one of the prominent models of stereotype change (Hewstone, 1989; Hewstone et al., 1992; Johnston & Hewstone, 1992; Weber & Crocker, 1983). What is the nature of that change?

If stereotyping rests on categorization, as argued earlier, and if category-based beliefs are applied to individuals as they are categorized, then the use of subtypes certainly alters the nature of the stereotyping that occurs. That is, if an individual is viewed as a member of a subtype, rather than being categorized as a member of the superordinate group, then this can have meaningful consequences. First, the beliefs associated with several of the subtypes may differ dramatically from those ascribed to the superordinate group (Brewer et al., 1981; Devine & Baker, 1991). If so, then the attributes assumed to characterize that individual could be substantially different from those inferred from the overall group stereotype, as well as from alternative subtypes. Second, although the use of subtypes still would constitute stereotyping (i.e., ascribing attributes on the basis of group membership), those attributes would be ascribed to a smaller, narrower range of persons (e.g., career women, urban Blacks, southern Democrats), rather than to the superordinate category as a whole (women, Blacks, Democrats). Thus, the overgeneralization that characterizes stereotyping could be reduced significantly (Hamilton, 1981).

These are important consequences that derive from the formation and use of subtypes. Nevertheless, there still will be occasions when the superordinate category will be considered an adequate basis for categorization and hence will be employed. The question then remains, does the presence of subtypes affect the nature of the beliefs ascribed to the overall group? Although this question has rarely been discussed in the stereotype literature, several possible answers can be gleaned from various discussions of subtyping.

One possibility is that the subtypes become sufficiently strong that they dilute the beliefs held about the superordinate category as a whole. For example, the Brewer et al. (1981) research on stereotypes of the elderly provided evidence that subjects held three different stereotypes about different subgroups of elderly persons. Each of these conceptions was richly articulated and widely shared in the subject population, yet the three were clearly different from each other. Moreover, these subjects had very little consensus about the attributes shared by old people in general (i.e., a stereotype that generalizes across the subtypes). Thus, the presence of well-developed (and presumably, frequently used) subtypes apparently has diluted the content of the stereotype of the generic category. To the extent that this happens, subtyping indeed would lead to stereotype change.

However, in other instances, this consequence is less evident. For example, Devine and Baker (1991) attempted to define and measure several subtypes of Blacks, including ghetto Blacks, Black athletes, streetwise Blacks, welfare Blacks, and businessman Blacks. They found that Black athletes and Black businessmen

formed distinct subtypes about whom subjects held clear, but differential, beliefs. On the other hand, ghetto Blacks, streetwise Blacks, and welfare Blacks were less distinct subgroups and shared overlapping attributes among themselves and with the generic stereotype of Blacks. In this case, it appears that certain subgroups (athletes, businessmen) are viewed as meaningful and unique subgroups of the overall category, whereas the stereotype of the superordinate category still applies, in large part, to most other group members (with other subgroups reflecting minor differences).

Finally, there may be cases in which highly specific, fairly narrow subtypes develop to represent disconfirming exceptions (Pettigrew, 1981) for the specific purpose of preserving the generic stereotype (Weber & Crocker, 1983). That is, if disconfirming cases can be separated into a distinct category that isolates them from the predominant group, their atypicality might diminish their perceived relevance for perceptions of the group as a whole. Thus, rather than producing stereotype change, the disconfirming instances become their own group, and thereby pose no challenge to the preexisting belief structure. In this case, subtyping becomes a mechanism contributing to stereotype preservation, rather than stereotype change. The implications of such "re-fencing" (Allport, 1954) are developed further in the next subsection.

C. *Nature of Disconfirming Instances.* The pattern in which disconfirming information is distributed among group members is only one of several properties of information that can affect the extent or likelihood of stereotype change. Another important factor concerns the extent to which a stereotype-disconfirming individual is perceived as being a "good" member of the social category about which the stereotype is held. As Rothbart and John (1985) pointed out, an individual person can be categorized in terms of any number of groups. They argued that if a person strongly violates the perceiver's conception of a particular group to which the target person belongs, the person will not be viewed as a member of that category and instead will be categorized in terms of some alternative group membership. For example, the football player who earns straight As, prefers fine wine, classical music, and English literature to beer, hard rock, and comic books, and shows warmth and sensitivity in his interpersonal relationships will not be categorized as a "jock," but rather as a "sophisticate" who happens to play a sport. That is, because he so completely violates the "jock" stereotype, he is not perceived as a member of that category. Consequently, the stereotype-disconfirming information that he provides is unlikely to alter the perceiver's stereotype of jocks; because he is not categorized as a jock, this information is not viewed as relevant to that stereotype. (Instead, his interest in playing football may produce some loosening of the perceiver's stereotype of "sophisticates," the social category that has been invoked by this individual.) Thus, there is a delicate balance between the extent to which an individual provides stereotype-disconfirming information and the ability of that information to have impact on the preexisting stereotype.

Rothbart and John (1985) argued that, to be an effective implement of change, the individual must (a) provide some stereotype-disconfirming information but (b) be perceived as a "good" member of the target group.

Experimental evidence supporting this argument was offered by Rothbart and Lewis (1988). They presented subjects with a number of exemplars of four geometric shapes (rectangles, triangles, etc.). For each shape, half of the exemplars were "poor" instances of the category (e.g., a rectangle where height is considerably longer than width), whereas the other half were "good" instances of the category (e.g., an equilateral triangle). For each shape, the "good" and "poor" exemplars were of different color, so that color was correlated perfectly with goodness of fit. Subsequently, subjects were asked to estimate how many instances of each shape they had seen in each color. Subjects estimated that they had seen significantly more of the "good" than of the "poor" examples of each category.

These findings have important implications for efforts to change stereotypes. Presenting disconfirming information will not be enough. In fact, presenting examples of individuals who strongly disconfirm stereotypic expectancies can be expected to have little, if any, effect, because those individuals will not be viewed as category members. Instead, group members who are in many respects representative of the target group, but who provide some disconfirmation of the stereotype, may be more effective in inducing change in stereotypic beliefs.

d. *Nature of Stereotypic Attributes.* Another factor that can influence the likelihood of stereotype change is the nature of the attributes comprising the stereotype (Hewstone, 1989). As Rothbart and Park (1986) showed, traits vary considerably in their susceptibility to change. Trait concepts are abstractions based on behavioral manifestations. For some traits, these behavioral manifestations are clearly evident and frequently observable, whereas for other traits those manifestations occur less frequently and are less apparent. For example, a person's cleanliness or dirtiness is readily observable and is apparent whenever that person is encountered; therefore, it is fairly easy for the perceiver to assess the validity of such a characterization and, if appropriate, to change one's belief about the person. On the other hand, for many other traits, such opportunities for validation are less frequent and are more difficult, rendering change in such a belief less likely. For example, consider the difficulty of disconfirming a belief that a person is devious. Behavioral manifestations that would provide clear evidence contradictory to this belief are difficult to imagine, would occur only under unusual (infrequent) circumstances, and often would be open to alternative interpretations that might undermine their potential for disconfirmation. In fact, even the absence of confirmatory manifestations can be viewed as substantiating the belief: the less one sees evidence of the person's deviousness, the more effective he or she appears to be at it. To the extent that a group's stereotype is composed of such traits, the more difficult it will be to induce change in those beliefs.

Similar considerations derive from Reeder and Brewer's (1979) analyses of the asymmetries in inferences from behavior to trait. For example, intelligent people can do stupid things, as well as smart things, but if a person is stupid, he or she does not have the capacity to do highly intelligent things. Similarly, dishonest people can and do engage in honest, as well as dishonest, behaviors. However, the honest person does not have such flexibility; as soon as he or she engages in a dishonest act, he or she is no longer an honest person. These asymmetries have implications for the likelihood of stereotype change. For instance, if a stereotype characterizes a group as dishonest, then honest behaviors by group members will not necessarily be viewed as compelling evidence against that belief, because dishonest persons are not always dishonest and, in fact, often engage in honest behaviors.

Finally, traits also differ in their breadth or narrowness (Hampson, John, & Goldberg, 1986). For example, the trait "responsible" is a broad trait that encompasses many diverse behaviors that might manifest this quality. In contrast, the trait "punctual" also refers to behavior that is responsible, but it represents a narrower, more restrictive, and more constrained domain of behavioral manifestations. The effects of these differences for stereotyping were demonstrated by Hamilton, Gibbons, Stroessner, and Sherman (1992), who analyzed subjects' ratings of liked and disliked nationalities on broad and narrow traits that differed in desirability. Not surprisingly, liked groups were rated highly on desirable traits and disliked groups were rated highly on undesirable traits. More interestingly, when subjects rated the groups on traits that were incompatible with their overall evaluations of the groups—rating liked groups on undesirable traits, or disliked groups on desirable traits—they made higher ratings on narrow than on broad attributes. That is, subjects acknowledged that a liked nationality possessed some undesirable attributes, and that disliked groups did have some favorable qualities, but these exceptions were confined to narrow traits that would pose less challenge to the preexisting group evaluation.

Although the research of Rothbart and Park (1986), Reeder (1985), and Hamilton et al. (1992) was not focused on issues of stereotype change directly, it points to important properties of trait-based beliefs that would make them more or less amenable to change. Given that group stereotypes are, to a considerable extent, composed of such trait-based beliefs, these results point to important considerations for an analysis of stereotype change, as well as indicating important avenues of future research on this topic.

e. *Outcome Bias and Stereotype Change.* It is obvious that stereotypes are difficult to change, and, in discussing this literature, we have encountered numerous reasons for their resistance to change even in the face of counterstereotypic information. From the cognitive perspective, the mechanisms that promote the use of stereotypes can undermine the effectiveness of efforts to change stereotypic beliefs. These mechanisms include processes and biases that, in many

circumstances, are functional and adaptive, even if they reflect less than optimal information processing. By and large, the cognitive processes we discussed serve to maintain the beliefs that generated, or biased, that processing from the outset. Recently, however, Mackie and her colleagues (Mackie, Allison, Worth, & Asuncion, 1992a, 1992b; Mackie, Worth, & Allison, 1990) have pursued a research program exploring the ways in which biases in information processing might be used to promote the perception that groups have changed over time, thereby facilitating the modification of stereotypic beliefs.

One of the best documented findings in the social perception literature is the correspondence bias—the tendency to infer that persons possess dispositions that correspond to the manifest properties of their behaviors, even when this behavior is clearly constrained by situational factors (Jones, 1991). This phenomenon also has been demonstrated for group perceptions (Allison & Messick, 1985). The Mackie et al. research capitalized on the fact that a perceiver's correspondent inferences can be biased by the outcome of the observed performance. For example, Mackie et al. (1990) showed that when a group of college students was described as having qualified for a "college bowl" competition, they were rated as being more intelligent than when they failed to qualify—a simple correspondent inference. However, this effect occurred (a) even when the group's success or failure was determined by arbitrary changes in the decision rules governing the competition, and (b) even when the group's actual behavioral performance was identical in the success and failure conditions.

More recently, this research has been extended in meaningful ways. First, Mackie et al. (1992b) showed that this outcome bias can generate biased inferences about groups even when those inferences are counter-stereotypic. Specifically, these effects occurred even when Asian-Americans were said to fail and Blacks were said to succeed in an intelligence-related competition (outcomes that are counter to the stereotypes held in the subject population). Finally, Mackie et al. (1992a) showed that outcome-biased inferences based on the performance of group members can generalize to members of the stereotyped group as a whole. Specifically, subjects who had made outcome-biased inferences about the intelligence of a group of eight Black or Asian-American college students were later asked to assess the intelligence of Whites, Blacks, and Asians in general. The results showed a modest, but significant, generalization of the outcome-biased inferences to the target group as a whole.

The research by Mackie and her colleagues demonstrated that cognitive biases in the way group-descriptive information is processed (in this case, inferences biased by performance outcomes) can influence judgments of the group as a whole, including the perception of change. These findings raise the possibility that other biasing factors also may be utilized as a means of inducing change in perceptions of, and ultimately beliefs about, stereotyped groups.

[6] A FINAL COMMENT

In preparing this chapter, we have surveyed a sizeable literature addressing a variety of topics related to the nature and use of stereotypes. In doing so, we were impressed by a number of features of this work. First, the pure magnitude of this literature is impressive, particularly when one realizes that a considerable proportion of the research we discussed was reported within the last 15 years. Clearly, this has been a vigorous area of research, and clearly the study of stereotypes is one domain that has been influenced in important ways by developments in social cognition. Second, we were impressed by the diversity of topics relevant to stereotyping that have been investigated during this period. This is not a narrow, singularly focused research area. We believe our review reveals the multifaceted nature of this topic and of how it is being explored empirically. Third, we were impressed by the sophistication of this research area, both conceptually and empirically. Theoretical models have been developed, debated, and tested that are rather specific about the nature of cognitive representations or about the processing of different kinds of information. Also the experimental strategies and research technology used to test those theoretical ideas have evolved at a rapid pace. Fourth, we were impressed by the broadening of this conceptual approach to stereotyping. In contrast to a decade ago, it is no longer a purely cognitive analysis, but instead seeks to understand the role of affective and motivational processes as well. In all of these fundamental respects, research on stereotypes has made impressive progress.

Social scientists have been studying the nature and functioning of stereotypes for several decades. They have approached this topic from several conceptual orientations and have used a variety of research tools in conducting their inquiry. It is clear that we are not very far down the path toward achieving a full understanding of this topic. Nevertheless, the breadth, depth, and vitality of the research we have reviewed provide encouraging omens of continued progress. As it has in the past, the specific nature of theorizing and research will continue to change and evolve. The research we have reviewed provides a solid foundation on which future developments can build.

ACKNOWLEDGMENTS

Preparation of this chapter was supported by NIMH Grant MH 40058 to the first author. The second author was supported by a National Science Foundation graduate fellowship. The authors thank Diane M. Mackie, Steven J. Sherman, and Robert S. Wyer, Jr., for their comments on preliminary versions of this chapter.

REFERENCES

Allison, S. T., & Messick, D. M. (1985). The group attribution error. *Journal of Experimental Social Psychology, 21*, 563–579.

Allport, G. W. (1954). *The nature of prejudice*. Cambridge, MA: Addison-Wesley.

Andersen, S. M., & Cole, S. W. (1990). "Do I know you?" The role of significant others in general social perception. *Journal of Personality and Social Psychology, 59*, 384–399.

Andersen, S. M., & Klatzky, R. L. (1987). Traits and social stereotypes: Levels of categorization in person perception. *Journal of Personality and Social Psychology, 53*, 235–246.

Ashmore, R. D. (1981). Sex stereotypes and implicit personality theory. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 37–81). Hillsdale, NJ: Lawrence Erlbaum Associates.

Ashmore, R. D., & Del Boca, F. K. (1979). Sex stereotypes and implicit personality theory: Toward a cognitive-social psychological conception. *Sex Roles, 5*, 219–248.

Ashmore, R. D., & Del Boca, F. K. (1981). Conceptual approaches to stereotypes and stereotyping. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 1–35). Hillsdale, NJ: Lawrence Erlbaum Associates.

Banaji, M. R., & Greenwald, A. G. (in press). Implicit stereotyping and prejudice. In M. P. Zanna & J. M. Olson (Eds.), *The psychology of prejudice, the Ontario Symposium* (Vol. 7). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bargh, J. A., Bond, R. N., Lombardi, W. L., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology, 50*, 869–879.

Bargh, J. A., & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology, 43*, 437–449.

Barsalou, L. W. (1990). On the indistinguishability of exemplar memory and abstraction in category representation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 3, pp. 61–88). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bechtold, A., Naccarato, M. E., & Zanna, M. P. (1986). *Need for structure and the prejudice-discrimination link*. Paper presented at the annual meeting of the Canadian Psychological Association, Toronto, Ontario.

Biernat, M., Manis, M., & Nelson, T. F. (1991). Stereotypes and standards of judgment. *Journal of Personality and Social Psychology, 60*, 485–499.

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

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. (1990). Stereotypes as judgmental heuristics: Evidence of circadian variations in discrimination. *Psychological Science, 1*, 319–322.

Bodenhausen, G. V. (1993). Emotion, arousal, and stereotypic judgments: A heuristic model of affect and stereotyping. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 13–37). San Diego, CA: Academic Press.

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., & Wyer, R. S., Jr. (1985). Effects of stereotypes on decision making and information processing strategies. *Journal of Personality and Social Psychology, 48*, 267–282.

Bower, G. H., & Karlin, M. B. (1974). Depth of processing pictures of faces and recognition memory. *Journal of Experimental Psychology, 4*, 751–757.

Brewer, M. B. (1979). In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological Bulletin, 86*, 307–324.

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: Lawrence Erlbaum Associates.

Brewer, M. B. (1993). Social identity, distinctiveness, and in-group homogeneity. *Social Cognition, 11*, 150–164.

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: Lawrence Erlbaum Associates.

Bruner, J. S. (1957). On perceptual readiness. *Psychological Review, 64*, 123–152.

Campbell, D. T. (1965). Ethnocentric and other altruistic motives. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 283–311). Lincoln, NE: University of Nebraska 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.

Chapman, L. J. (1967). Illusory correlation in observational report. *Journal of Verbal Learning and Verbal Behavior, 6*, 151–155.

Clary, E. G., & Tesser, A. (1983). Reactions to unexpected events: The naive scientist and interpretive activity. *Personality and Social Psychology Bulletin, 9*, 609–620.

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.

Crocker, J., & Major, B. (1989). Social stigma and self-esteem: The self-protective properties of stigma. *Psychological Review, 96*, 608–630.

Darley, J. M., & Gross, P. H. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology, 44*, 20–33.

Deaux, K., Winton, W., Crowley, M., & Lewis, L. L. (1985). Level of categorization and content of gender stereotypes. *Social Cognition, 3*, 145–167.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5–18.

Devine, P. G., & Baker, S. M. (1991). Measurement of racial stereotype subtyping. *Personality and Social Psychology Bulletin, 17*, 44–50.

Devine, P. G., & Monteith, M. J. (1993). The role of discrepancy-associated affect in prejudice reduction. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 317–344). San Diego, CA: Academic Press.

Devine, P. G., Monteith, M. J., Zuwerink, J. R., & Elliot, A. J. (1991). Prejudice with and without compunction. *Journal of Personality and Social Psychology, 60*, 817–830.

Diehl, M. (1990). The minimal group paradigm: Theoretical explanations and empirical findings. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 1, pp. 263–292). Chichester, England: Wiley.

Dijker, A. G. M. (1987). Emotional reactions to ethnic minorities. *European Journal of Social Psychology, 17*, 305–326.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representation. *Journal of Experimental Social Psychology, 22*, 22–37.

Dovidio, J. F., & Gaertner, S. L. (1993). Stereotypes and evaluative intergroup bias. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 167–193). San Diego, CA: Academic Press.

Duncan, B. L. (1976). Differential social perception and attribution of intergroup violence: Testing the lower limits of stereotyping of blacks. *Journal of Personality and Social Psychology, 34*, 590–598.

Esses, V. M., Haddock, G., & Zanna, M. P. (1993). Values, stereotypes, and emotions as determinants of intergroup attitudes. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 137–166). San Diego, CA: Academic Press.

Fazio, R. H., Lenn, T. M., & Effrein, E. A. (1984). Spontaneous attitude formation. *Social Cognition, 2*, 217–234.

Fiedler, K. F. (1991). The tricky nature of skewed frequency tables: An information loss account of distinctiveness-based illusory correlations. *Journal of Personality and Social Psychology, 60*, 24–36.

Fiske, S. T. (1982). Schema-triggered affect: Applications to social perception. In M. S. Clark & S. T. Fiske (Eds.), *Affect and cognition* (pp. 55–78). Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

Fiske, S. T., Neuberg, S. L., Beattie, A. E., & Milberg, S. J. (1987). Category-based and attribute-based reactions to others: Some informational conditions of stereotyping and individuating processes. *Journal of Experimental Social Psychology, 23*, 399–427.

Fiske, S. T., & Pavelchak, M. A. (1986). Category-based versus piecemeal-based affective responses: Developments in schema-triggered affect. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 1, pp. 167–203). New York: Guilford.

Fiske, S. T., & Ruscher, J. B. (1993). Negative interdependence and prejudice: Whence the affect? In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 239–268). San Diego, CA: Academic Press.

Fiske, S. T., & Taylor, S. E. (1991). *Social cognition*. New York: McGraw-Hill.

Forgas, J. P. (1992). Affect and social perception: Research evidence and an integrative theory. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 3, pp. 183–223). Chichester, England: Wiley.

Forgas, J. P., & Moylan, S. J. (1991). Affective biases in stereotype judgments. *Cognition and Emotion, 5*, 379–395.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 234–257.

Gilbert, D. T., & Hixon, J. G. (1991). The trouble of thinking: Activation and application of stereotypic beliefs. *Journal of Personality and Social Psychology, 60*, 509–517.

Gilbert, G. M. (1951). Stereotype persistence and change among college students. *Journal of Abnormal and Social Psychology, 46*, 245–254.

Gilovich, T. (1981). Seeing the past in the present: The effect of associations to familiar events on judgments and decisions. *Journal of Personality and Social Psychology, 40*, 797–808.

Gurwitz, S. B., & Dodge, K. A. (1977). Effects of confirmations and disconfirmations on stereotype-based attributions. *Journal of Personality and Social Psychology, 35*, 495–500.

Hamilton, D. L. (1979). A cognitive-attributional analysis of stereotyping. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 12, pp. 53–84). New York: Academic Press.

Hamilton, D. L. (1981). Stereotyping and intergroup behavior: Some thoughts on the cognitive approach. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 333–353). Hillsdale, NJ: Lawrence Erlbaum Associates.

Hamilton, D. L., Driscoll, D. M., & Worth, L. T. (1989). Cognitive organization of impressions: Effects of incongruency in complex representations. *Journal of Personality and Social Psychology, 57*, 925–939.

Hamilton, D. L., Dugan, P. M., & Trolier, 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., Gibbons, P., Stroessner, S. J., & Sherman, J. W. (1992). Stereotypes and language use. In K. Fiedler & G. R. Semin (Eds.), *Language and social cognition* (pp. 102–128). Newbury Park, CA: Sage.

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

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: Lawrence Erlbaum Associates.

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, S. J. (1989). Illusory correlations: Implications 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., & Ruvolo, C. M. (1990). Stereotype-based expectancies: Effects on information processing and social behavior. *Journal of Social Issues, 46*(2), 35–60.

Hamilton, D. L., Stroessner, S. J., & Driscoll, D. M. (in press). Social cognition and the study of stereotypes. In P. G. Devine, D. L. Hamilton, & T. M. Ostrom (Eds.), *Social cognition: Contributions to classic issues in social psychology*. New York: Springer-Verlag.

Hamilton, D. L., Stroessner, S. J., & Mackie, D. M. (1993). The influence of affect on stereotyping: The case of illusory correlations. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 39–61). San Diego, CA: Academic Press.

Hamilton, D. L., & Trolier, 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.

Hampson, S. E., John, O. P., & Goldberg, L. R. (1986). Category breadth and hierarchical structure in personality: Studies in asymmetries in judgments of trait implications. *Journal of Personality and Social Psychology, 51*, 37–54.

Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology, 46*, 44–56.

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.

Hayes-Roth, B., & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning and Verbal Behavior, 16*, 321–338.

Heilman, M. E. (1984). Information as a deterrent against sex discrimination: The effects of applicant sex and information type on preliminary employment decisions. *Organizational Behavior and Human Performance, 33*, 174–186.

Hewstone, M. (1989). Changing stereotypes with disconfirming information. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotypes and prejudice: Changing conceptions* (pp. 207–223). New York: Springer-Verlag.

Hewstone, M., Johnston, L., & Aird, P. (1992). Cognitive models of stereotype change (2): Perceptions of homogeneous and heterogeneous groups. *European Journal of Social Psychology, 22*, 235–249.

Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). The nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 59–69.

Higgins, E. T., & King, G. A. (1981). Accessibility of social constructs: Information-processing consequences of individual and contextual variability. In N. Cantor & J. F. Kihlstrom (Eds.), *Personality, cognition, and social interaction* (pp. 69–122). Hillsdale, NJ: Lawrence Erlbaum Associates.

Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology, 13*, 141–154.

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.

Isen, A. M. (1984). Toward understanding the role of affect in cognition. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 3, pp. 179–236). Hillsdale, NJ: Lawrence Erlbaum Associates.

Jackson, L. A., & Sullivan, L. A. (1989). Cognition and affect in evaluations of stereotyped group members. *The Journal of Social Psychology, 129*, 659–672.

Jacoby, L. L., Kelley, C. M., Brown, J., & Jasechko, J. (1989). Becoming famous overnight: Limits on the ability to avoid unconscious influences of the past. *Journal of Personality and Social Psychology, 56*, 326–338.

Johnson, C., & Mullen, B. (in press). Evidence for the accessibility of paired distinctiveness in distinctiveness-based illusory correlation in stereotyping. *Personality and Social Psychology Bulletin.*

Johnston, L., & Hewstone, M. (1992). Cognitive models of stereotype change (3): Subtyping and the perceived typicality of disconfirming group members. *Journal of Experimental Social Psychology, 28*, 360–386.

Jones, E. E. (1991). *Interpersonal perception*. New York: W. H. Freeman.

Jones, E. E., Wood, G. C., & Quattrone, G. A. (1981). Perceived variability of personal characteristics in in-groups and out-groups: The role of knowledge and evaluation. *Personality and Social Psychology Bulletin, 7*, 523–528.

Jones, R. A., & Ashmore, R. D. (1973). The structure of intergroup perception: Categories and dimensions in views of ethnic groups and adjectives used in stereotype research. *Journal of Personality and Social Psychology, 25*, 428–438.

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.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review, 93*, 136–153.

Karasawa, M., & Brewer, M. B. (1992). *Category size and judgments of variability: The effects of seeing the trees in the forest*. Unpublished manuscript, University of California, Los Angeles.

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.

Kashima, E. S., & Kashima, Y. (1993). Perceptions of general variability of social groups. *Social Cognition, 11*, 1–21.

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

Katz, I. (1981). *Stigma: A social psychological analysis*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Katz, P. A. (1976). The acquisition of racial attitudes in children. In P. A. Katz (Ed.), *Towards the elimination of racism* (pp. 125–154). New York: Pergamon.

Kim, H-S., & Baron, R. S. (1988). Exercise and the illusory correlation: Does arousal heighten stereotypic processing? *Journal of Experimental Social Psychology, 24*, 366–380.

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.

Kraus, S., Ryan, C. S., Judd, C. M., Hastie, R., & Park, B. (1993). Use of mental frequency distributions to represent variability among members of social categories. *Social Cognition, 11*, 22–43.

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.

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 of 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.

Linville, P. W., Salovey, P., & Fischer, G. W. (1986). Stereotyping and perceived distributions of social characteristics: An application to in-group–out-group perception. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 165–208). Orlando, FL: Academic Press.

Lippmann, W. (1922). *Public opinion*. New York: Harcourt, Brace, Jovanovitch.

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

Lord, C. G., & Saenz, D. S. (1985). Memory deficits and memory surfeits: Differential cognitive consequences of tokenism for tokens and observers. *Journal of Personality and Social Psychology, 49*, 918–926.

Maass, A., & Schaller, M. (1991). Intergroup biases and the cognitive dynamics of stereotype formation. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 2, pp. 189–209). Chichester, England: Wiley.

Mackie, D. M., Allison, S. T., Worth, L. T., & Asuncion, A. G. (1992a). The generalization of outcome-biased counter-stereotypic inferences. *Journal of Experimental Social Psychology, 28*, 43–64.

Mackie, D. M., Allison, S. T., Worth, L. T., & Asuncion, A. G. (1992b). The impact of outcome biases on counterstereotypic inferences about groups. *Personality and Social Psychology Bulletin, 18*, 44–51.

Mackie, D. M., Asuncion, A. G., & Rosselli, F. (1992). The impact of positive affect on persuasion processes. In M. S. Clark (Ed.), *Review of personality and social psychology* (Vol. 14, pp. 247–270). Newbury Park, CA: Sage.

Mackie, D. M., & Hamilton, D. L. (1993). *Affect, cognition, and stereotyping: Interactive processes in group perception*. San Diego, CA: Academic Press.

Mackie, D. M., Hamilton, D. L., Schroth, H. A., Carlisle, C. J., Gersho, B. F., Meneses, L. M., Nedler, B. F., & Reichel, L. D. (1989). The effects of induced mood on illusory correlations. *Journal of Experimental Social Psychology, 25*, 524–544.

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.

Mackie, D. M., & Worth, L. T. (1991). "Feeling good but not thinking straight": Positive mood and persuasion. In J. P. Forgas (Ed.), *Emotion and social judgments* (pp. 201–220). Oxford: Pergamon.

Mackie, D. M., Worth, L. T., & Allison, S. T. (1990). Outcome biased inferences and the perception of change in groups. *Social Cognition, 8*, 325–342.

Macrae, C. N., Hewstone, M., & Griffiths, R. J. (1993). Processing load and memory for stereotype-based information. *European Journal of Social Psychology, 23*, 77–87.

Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (in press). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of Personality and Social Psychology.*

Macrae, C. N., & Shepherd, J. W. (1989). Stereotypes and social judgments. *British Journal of Social Psychology, 28*, 319–325.

Macrae, C. N., Shepherd, J. W., & Milne, A. B. (1992). The effects of source credibility on the dilution of stereotype-based judgments. *Personality and Social Psychology Bulletin, 18*, 765–775.

Major, B., & Crocker, J. (1993). Social stigma: The consequences of attributional ambiguity. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 345–370). San Diego, CA: Academic Press.

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.

Medin, D. L., Altom, M. W., & Murphy, T. D. (1984). Given versus induced category representations: Use of prototype and exemplar information in classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 333–352.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207–238.

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.

Messick, D. M., & Mackie, D. M. (1989). Intergroup relations. In M. R. Rosenzweig & L. W. Porter (Eds.), *Annual review of psychology* (Vol. 40, pp. 45–81). Palo Alto, CA: Annual Reviews.

Miller, N., & Brewer, M. B. (1986). Categorization effects on in-group and out-group perception. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 209–230). Orlando, FL: Academic Press.

Mullen, B., & Hu, L. (1989). Perceptions of in-group and out-group variability: A meta-analytic integration. *Basic and Applied Social Psychology, 10*, 233–252.

Mullen, B., & Johnson, C. (1990). Distinctiveness-based illusory correlations and stereotyping: A meta-analytic integration. *British Journal of Social Psychology, 29*, 11–28.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review, 92*, 289–316.

Neuberg, S. L., & Fiske, S. T. (1987). Motivational influences on impression formation: Outcome dependency, accuracy-driven attention, and individuating processes. *Journal of Personality and Social Psychology, 53*, 431–444.

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.

Oakes, P. J., & Turner, J. C. (1990). Is limited information processing capacity the cause of social stereotyping? In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 1, pp. 111–135). Chichester, England: Wiley.

Ostrom, T. M., & Sedikides, C. (1992). Out-group homogeneity effects in natural and minimal groups. *Psychological Bulletin, 112*, 536–552.

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. (1990). Measures and models of perceived group variability. *Journal of Personality and Social Psychology, 59*, 173–191.

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, 211–245). Chichester, England: Wiley.

Park, B., & Rothbart, M. (1982). Perception of out-group homogeneity and levels of social categorization: Memory for the subordinate attributes of in-group and out-group members. *Journal of Personality and Social Psychology, 42*, 1051–1068.

Park, B., Ryan, C. S., & Judd, C. M. (1992). The role of meaningful subgroups in explaining differences in perceived variability for in-groups and out-groups. *Journal of Personality and Social Psychology, 63*, 553–567.

Pettigrew, T. F. (1981). Extending the stereotype concept. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 303–332). Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

Quattrone, G. A. (1986). On the perception of a group's variability. In S. Worchel & W. Austin (Eds.), *The psychology of intergroup relations* (2nd ed., pp. 25–48). Chicago: Nelson-Hall.

Quattrone, G. A., & Jones, E. E. (1980). The perception of variability within in-groups and out-groups: Implications for the Law of Small Numbers. *Journal of Personality and Social Psychology, 38*, 141–152.

Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86*, 61–79.

Rojahn, K., & Pettigrew, T. F. (1992). Memory for schema-relevant information: A meta-analytic resolution. *British Journal of Social Psychology, 31*, 81–110.

Rosch, E. H. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 27–48). Hillsdale, NJ: Lawrence Erlbaum Associates.

Rothbart, M. (1981). Memory processes and social beliefs. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 145–182). Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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.

Rothbart, M., & John, O. P. (1985). Social categorization and behavioral episodes: A cognitive analysis of the effects of intergroup contact. *Journal of Social Issues, 41*(3), 81–104.

Rothbart, M., & Lewis, S. (1988). Inferring category attributes from exemplar attributes: Geometric shapes and social categories. *Journal of Personality and Social Psychology, 55*, 861–872.

Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology, 50*, 131–142.

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.

Sanbonmatsu, D. M., Shavitt, S., Sherman, S. J., & Roskos-Ewoldsen, D. R. (1987). Illusory correlation in the perception of performance by self or a salient other. *Journal of Experimental Social Psychology, 23*, 518–543.

Sanbonmatsu, D. M., Sherman, S. J., & Hamilton, D. L. (1987). Illusory correlation in the perception of individuals and groups. *Social Cognition, 5*, 1–25.

Schaller, M. (1991). Social categorization and the formation of group stereotypes: Further evidence for biased information processing in the perception of group-behavior correlations. *European Journal of Social Psychology, 21*, 25–35.

Schaller, M., & Maass, A. (1989). Illusory correlation and social categorization: Toward an integration of motivational and cognitive factors in stereotype formation. *Journal of Personality and Social Psychology, 56*, 709–721.

Schwarz, N. (1990). Feelings as information: Informational and motivational functions of affective states. In R. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 527–561). New York: Guilford.

Sherif, M. (1967). *Group conflict and cooperation*. London: Routledge & Kegan Paul.

Sherman, J. W., & Hamilton, D. L. (in press). On the formation of interitem links in person memory. *Journal of Experimental Social Psychology*.

Simon, B., & Brown, R. (1987). Perceived intragroup homogeneity in minority–majority contexts. *Journal of Personality and Social Psychology, 53*, 703–711.

Simon, B., & Hamilton, D. L. (in press). Self-stereotyping and social context: The effects of relative in-group size and in-group status. *Journal of Personality and Social Psychology*.

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.

Slusher, M. P., & Anderson, C. A. (1987). When reality monitoring fails: The role of imagination in stereotype maintenance. *Journal of Personality and Social Psychology, 52*, 653–662.

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: Lawrence Erlbaum Associates.

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

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.

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.

Srull, T. K., Lichtenstein, M., & Rothbart, M. (1985). Associative storage and retrieval processes in person memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 11*, 316–345.

Srull, T. K., & Wyer, R. S., Jr. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology, 37*, 1660–1672.

Srull, T. K., & Wyer, R. S., Jr. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology, 38*, 841–856.

Srull, T. K., & Wyer, R. S., Jr. (1989). Person memory and judgment. *Psychological Review, 96*, 58–83.

Stangor, C., & Duan, C. (1991). Effects of multiple task demands upon memory for information about social groups. *Journal of Experimental Social Psychology, 27*, 357–378.

Stangor, C., & Ford, T. E. (1992). Accuracy and expectancy-confirming processing orientations and the development of stereotypes and prejudice. In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 3, pp. 57–89). Chichester, England: Wiley.

Stangor, C., & McMillan, D. (1992). Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychological Bulletin, 111*, 42–61.

Stangor, C., & Ruble, D. N. (1989). Stereotype development and memory: What we remember depends on how much we know. *Journal of Experimental Social Psychology, 25*, 18–35.

Stangor, C., Sullivan, L. A., & Ford, T. E. (1991). Affective and cognitive determinants of prejudice. *Social Cognition, 9*, 359–380.

Stephan, W. G. (1985). Intergroup relations. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology* (3rd ed., Vol. 2, pp. 599–658). New York: Random House.

Stephan, W. G., & Stephan, C. W. (1985). Intergroup anxiety. *Journal of Social Issues, 41*(3), 157–175.

Stephan, W. G., & Stephan, C. W. (1993). Cognition and affect in stereotyping: Parallel interactive networks. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 111–136). San Diego, CA: Academic Press.

Stern, L. D., Marrs, S., Millar, M. G., & Cole, E. (1984). Processing time and the recall of inconsistent and consistent behaviors of individuals and groups. *Journal of Personality and Social Psychology, 47*, 253–262.

Stroebe, W., & Insko, C. A. (1989). Stereotyping, prejudice, and discrimination: Changing conceptions in theory and research. In D. Bar-Tal, C. F. Graumann, A. W. Kruglanski, & W. Stroebe (Eds.), *Stereotypes and prejudice: Changing conceptions* (pp. 3–36). New York: Springer-Verlag.

Stroessner, S. J. (1992). *Target-based determinants of social categorization: Extensions of the "white male norm" hypothesis*. Unpublished doctoral dissertation, University of California, Santa Barbara.

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.

Stroessner, S. J., & Mackie, D. M. (1992). The impact of induced affect on the perception of variability in social groups. *Personality and Social Psychology Bulletin, 18*, 546–554.

Stroessner, S. J., & Mackie, D. M. (1993). Affect and perceived group variability: Implications for stereotyping and prejudice. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 63–86). San Diego, CA: Academic Press.

Swann, W. B., Jr. (1984). Quest for accuracy in person perception: A matter of pragmatics. *Psychological Review, 91*, 457–477.

Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of Social Issues, 25*(4), 79–97.

Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American, 223*(5), 96–102.

Tajfel, H., Billig, M., Bundy, R. P., & Flament, C. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology, 1*, 149–177.

Tajfel, H., & Turner, J. C. (1979). An integrative theory of intergroup conflict. In W. G. Austin & S. Worchel (Eds.), *The social psychology of intergroup relations* (pp. 33–47). Monterey, CA: Brooks/Cole.

Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. Austin (Eds.), *The psychology of intergroup relations* (2nd ed., pp. 7–24). Chicago: Nelson-Hall.

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

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: Lawrence Erlbaum Associates.

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: Lawrence Erlbaum Associates.

Taylor, S. E., Fiske, S. T., Etcoff, N. L., & Ruderman, A. (1978). Categorical bases of person memory and stereotyping. *Journal of Personality and Social Psychology, 36*, 778–793.

Turner, J. C. (1987). *Rediscovering the social group: A self-categorization theory*. New York: Basil Blackwell.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327–352.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207–232.

Vanbeselaere, N. (1991). The different effects of simple and crossed categorizations: A result of the category differentiation process or of differential category salience? In W. Stroebe & M. Hewstone (Eds.), *European review of social psychology* (Vol. 2, pp. 247–278). Chichester, England: Wiley.

Vanman, E. J., & Miller, N. (1993). Applications of emotion theory and research to stereotyping and intergroup relations. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 213–238). San Diego, CA: Academic Press.

Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology, 45*, 961–977.

Wilder, D. A. (1978). Homogeneity of jurors: The majority's influence depends upon their perceived independence. *Law and Human Behavior, 2*, 363–376.

Wilder, D. A. (1981). Perceiving persons as a group: Categorization and intergroup relations. In D. L. Hamilton (Ed.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 213–258). Hillsdale, NJ: Lawrence Erlbaum Associates.

Wilder, D. A. (1993). The role of anxiety in facilitating stereotypic judgment of out-group behavior. In D. M. Mackie & D. L. Hamilton (Eds.), *Affect, cognition, and stereotyping: Interactive processes in group perception* (pp. 87–109). San Diego, CA: Academic Press.

Wilder, D. A., & Shapiro, P. (1989). Role of competition-induced anxiety in limiting the beneficial impact of positive behavior by an out-group member. *Journal of Personality and Social Psychology, 56*, 60–69.

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.

Zadny, J., & Gerard, H. B. (1974). Attributed intentions and informational selectivity. *Journal of Experimental Social Psychology, 10*, 34–52.

Zarate, M. A., & Smith, E. R. (1990). Person categorization and stereotyping. *Social Cognition, 8*, 161–185.

2 Cognitive Processes in Attitude Change

Richard E. Petty
Joseph R. Priester
Duane T. Wegener
The Ohio State University

Contents

THE NATURE AND PROPERTIES OF ATTITUDES AND ATTITUDE SYSTEMS

In the 1992 presidential election, public opinion polls showed large swings in people's attitudes toward the candidates over relatively short time periods. About a year before the actual election, the incumbent President George Bush was held in high esteem by about 80% of the electorate. By the time of the Democratic National Convention, his approval rating plunged to around 30%, and his