

Forming Impressions From Stereotypes, Traits, and Behaviors: A Parallel-Constraint-Satisfaction Theory

Ziva Kunda and Paul Thagard
University of Waterloo

The authors propose a parallel-constraint-satisfaction theory of impression formation that assumes that social stereotypes and individuating information such as traits or behaviors constrain each other's meaning and jointly influence impressions of individuals. Building on models of text comprehension (W. Kintsch, 1988), the authors describe a connectionist model that can account for the major findings on how stereotypes affect impressions of individuals in the presence of different kinds of individuating information; how stereotypes, behaviors, and traits affect each other's meaning; and how multiple stereotypes jointly affect impressions. Most of these findings can be modeled by constraint networks, which suggests that they may be due to relatively automatic processes that require little conscious inference. The authors also point to a small number of phenomena that involve more controlled processes. The advantages of the authors' parallel model over serial models are discussed.

How friendly is the person you have just met at the party? How aggressive is your neighbor? How decisive is President Clinton? Your answers to such questions may be shaped by your understanding of the implications of these people's behavior, profession, age, sex, ethnicity, interpersonal relations, personality traits, physical appearance, abilities, goals, family background, or any other information about them that you consider relevant. These diverse varieties of information that can shape our impressions of others are typically classified by social psychologists into two major kinds—stereotypes and individuating information (Brewer, 1988; Fiske & Neuberg, 1990; Locksley, Hepburn, & Ortiz, 1982). Stereotypes refer to membership in social categories such as sex, race, age, or profession that are believed to be associated with certain traits and behaviors. Individuating information refers to anything else known about the individual—behavior (e.g., hit someone), personality (e.g., introverted), family circumstances (e.g., has two brothers), etc. As we shall discuss later, the distinction between these two kinds of information is not always clear. But many theorists have treated them as qualitatively different, and a great deal of theoretical and empirical effort has gone into determining how an

impression is formed of an individual by combining the applicable stereotypes and individuating information.

Early theorists, most notably Asch (1946) and N. H. Anderson (1968), assumed that the full range of information known to characterize an individual is integrated into one's impression of that individual. Asch and Anderson differed in their accounts of how this integration was achieved. From a gestalt perspective, Asch (1946; 1952/1987; Asch & Zukier, 1984) assumed that the perceiver blends the diverse features of the target person into a coherent, unitary impression that takes into account the meaning of individual features as well as their interrelationships. In the process, the various aspects of the person are modified and made to fit together coherently. N. H. Anderson (1968), taking an elementaristic perspective, assumed that the perceiver assesses the implications of each piece of information about the target person separately and then combines them algebraically into a summary impression. Despite their different accounts of the process of information integration, both these approaches share the view that one's impression of a person reflects an integration of all the information known to characterize that person.

This view has been challenged by two more recent models, Brewer's (1988) *dual process model* and Fiske and Neuberg's (1990) *continuum model* of impression formation.¹ Although the two models differ from each other in important ways (Fiske, 1990), they share some key assumptions. Both assume that it is necessary to distinguish between stereotype- or category-based, top-down processes on the one hand, and attribute-based, piecemeal, bottom-up processes on the other hand. Brewer assumes

Ziva Kunda, Psychology Department, University of Waterloo, Waterloo, Ontario, Canada; Paul Thagard, Philosophy Department, University of Waterloo.

This research was supported by grants by the Social Sciences and Humanities Research Council of Canada and by the Natural Sciences and Engineering Research Council of Canada.

We are grateful to John Holmes, Walter Kintsch, Michael Ross, Stephen Read, and Mark Zanna for comments on an earlier version of this article, and to Jeff Poupore for help with programming.

Correspondence concerning this article should be addressed to Ziva Kunda, Psychology Department, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1. Electronic mail may be sent via Internet to zkunda@watarts.uwaterloo.ca.

¹ In describing Brewer's (1988) and Fiske and Neuberg's (1990) models, we focus only on their implications for the question of how stereotypes are integrated with individuating information. Both models also address other important issues, such as category activation and the role of attention, that we do not discuss in detail.

that perceivers engage in either one or the other of these processes, whereas Fiske and Neuberg assume that the two represent the extreme endpoints of a continuum that also has intermediate kinds of processes. Both assume that the different types of information—stereotypes and individuating information—will contribute to impression formation differently, depending on the kind of process engaged in. Impressions resulting from stereotype-based processes will be dominated by the stereotype to the point that attributes and their implications may be ignored, whereas impressions resulting from attribute-based processes will be dominated by the attributes. Both models are serial in that they assume that individuals first, perhaps automatically, engage in stereotype-based processes. They may then engage in more attribute-based processes if they are strongly motivated to (Brewer, 1988; Fiske & Neuberg, 1990) or if the nature of the attributes prevents attempts at categorizing or recategorizing the individual as belonging to any particular stereotype (Fiske & Neuberg, 1990). Thus both models assume that stereotype-based processes dominate the more attribute-based ones.

Building on recent developments in theorizing about cognitive processes, we propose an alternative, parallel-constraint-satisfaction account of impression formation. We propose that rather than being processed serially, stereotypes and individuating information are processed simultaneously, and jointly influence impressions of individuals as well as each other's meaning. In proposing this parallel processing model, we return full circle to earlier views of person perception, especially to Asch's gestalt approach. Indeed, our proposal may be viewed as a more precise implementation of gestalt ideas (Read & Miller, 1993; Spellman & Holyoak, 1992).

In the last decade, an important new trend has emerged in cognitive science. Rather than viewing cognitive processes as serial, many models of cognition now view them as occurring in parallel and as simultaneously constraining each other (e.g., Holyoak & Spellman, 1993; Rumelhart & McClelland, 1986). Parallel-constraint-satisfaction models have been applied to many perceptual and cognitive processes. These include low-level processes such as those involved in stereoscopic vision (Marr & Poggio, 1976) and letter perception (McClelland & Rumelhart, 1981) as well as higher level cognitive processes such as discourse comprehension (Kintsch, 1988), explanation evaluation (Thagard, 1989, 1992), analogical mapping (Holyoak & Thagard, 1989, 1995), and decision making (Thagard & Millgram, 1995). Applications have also been made to social-psychological processes such as dissonance reduction (Read & Miller, 1994; Shultz & Lepper, 1992), trait inference (Read & Miller, 1993), and social explanation (Read & Marcus-Newhall, 1993) as well as to understanding the nature of personality traits (Mischel & Shoda, 1995). Our parallel-constraint-satisfaction model of impression formation capitalizes on these recent developments.

After describing our theory and model, we review the literature on how stereotypes and individuating information affect impressions of individuals and on how they affect each other's meaning. We characterize the findings emerging from this literature as a series of phenomena. In so doing, we specify the circumstances under which stereotypes do and do not affect vari-

ous aspects of perceivers' impressions of individuals and describe the kinds of effects that stereotypes and individuating information can have on each other's meaning. We show how our parallel-constraint-satisfaction model of impression formation can account for each of these phenomena. Our model accounts for all phenomena that the earlier, serial models proposed by Brewer (1988) and by Fiske and Neuberg (1990) can account for, as well as for several phenomena that the serial models cannot readily explain.

To substantiate our account, we have implemented our parallel-constraint-satisfaction model in a connectionist computer program. We describe how this program simulates many important experimental findings about impression formation. Developing the program required us to spell out our assumptions in far greater detail than has been done by any previous theory of impression formation. Successful simulations of experimental results with this program suggest that the assumptions embodied in our model are computationally feasible and psychologically plausible.

Impression Formation as Parallel Constraint Satisfaction

We view the social perceiver's task of trying to make sense of incoming information about a person as similar to the task of text comprehension (Galambos, Abelson, & Black, 1986; Kintsch, 1988; Read & Miller, 1993). The social perceiver needs to decipher and integrate the meanings of incoming pieces of information about a target person and does so on the basis of a preexisting knowledge base that includes representations of social constructs such as stereotypes, traits, and behaviors as well as the interrelationships among these constructs. We want to model a perceiver who has noticed several types of information about a target person and forms an impression of the person by integrating this information. We do not address the processes that govern which types of information will be attended to in the first place (for a detailed discussion of the determinants and consequences of attention allocation, see Fiske & Neuberg, 1990). Most of the empirical research on how stereotypes are combined with individuating information to form an impression of an individual assumes that perceivers are aware of all of the information provided to them; the question of interest is how they make use of it.

Theory

We introduce our theory by analyzing the classic demonstration of the effects of stereotypes on the interpretation of the behavior of a person observed elbowing another person. It has been shown that this behavior is interpreted as a jovial shove when performed by a White person, but as a violent push when performed by a Black person (Duncan, 1976; Sagar & Schofield, 1980).

Our theory of impression formation assumes that stereotypes, traits, and behaviors can be represented as interconnected nodes in a spreading activation network. The spread of activation between nodes is constrained by positive and negative associations. Unlike earlier spreading activation models (J. An-

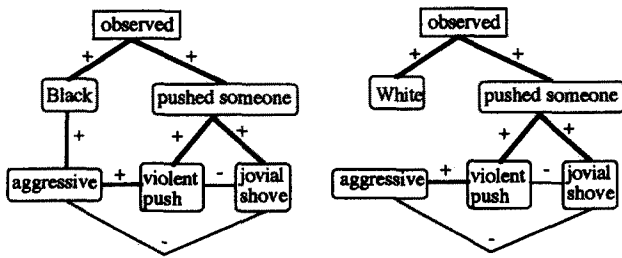


Figure 1. Stereotypes affect the meaning of behavior (Phenomenon 1). *Violent push* is activated more strongly in the network on the left than in the network on the right.

person, 1983; Collins & Loftus, 1975; Quillian, 1968), which assumed that all associations are positive and so activation may spread indefinitely, we assume that nodes can not only activate each other but also deactivate each other.

Figure 1 portrays a schematic illustration of a segment of the knowledge net that would come into play when a Black or a White person is observed pushing someone. In this and all the following figures, boxes depict the nodes representing the behavior (pushed someone), some of its possible interpretations (violent push and jovial shove), and the stereotyped categories (Black and White). The lines connecting these constructs indicate the nature of the connections among them. Bold lines indicate excitatory connections, and thin lines indicate inhibitory connections. The number of plusses or minuses alongside these lines indicates connection strength. Each of the constructs depicted also has many additional associates that are not portrayed in the figure. In this illustration, *aggressive* has a positive, excitatory link with *violent push* and a negative, inhibitory link with *jovial shove*. Therefore, when *aggressive* is activated, it will activate *violent push* and deactivate *jovial shove*. *Jovial shove* will be further deactivated because it also has a negative link with *violent push*, so that the more *violent push* is activated, the more *jovial shove* is deactivated.

Our theory further assumes that information that has been directly observed, that is, information known to be true of the target person, constrains the impression formed of that person. Figure 1 depicts the special status of observed information by connecting such information to a node termed *observed*. In this manner observed information may be distinguished from inferred knowledge. In this example the behavior (*pushed someone*) and the stereotyped category (*Black* or *White*) are observed. The possible interpretations of the behavior and the traits associated with the stereotype are not themselves observed but become activated or deactivated through their associations with the observed information. These associations reflect perceivers' preexisting beliefs.

Finally, we assume that impression formation occurs holistically by parallel constraint satisfaction performed by spreading activation. In other words, the associates of the observed information are activated and deactivated simultaneously and jointly constrain the impression of the target.

In our example, when one observes that a Black person pushed someone, *pushed someone* activates both *violent push* and *jovial shove*. At the same time, *Black* activates *aggressive*

which, in turn, activates *violent push* and deactivates *jovial shove*. In contrast, when the person is White, *aggressive* does not get activated by *White*. Therefore, *violent push* ends up with more activation when the same pushing behavior is performed by a Black than by a White. This is how stereotypes can constrain the meaning of behavior.

As illustrated by this example, the meaning of social constructs, as for that of any other concept, is not defined in the knowledge net. Rather, the meaning arises from the construct's location in the net. The meaning of a construct at a given moment may be conceived of as the subset of its associates that are activated at that moment. Ultimately, all nodes in the network are connected to all other ones, and the full meaning of a given construct may only be obtained through full knowledge of the entire network. However, the entire network is never activated as a whole. Rather, at any given moment only a small subset of a construct's associates are activated, and the nature of this subset depends on the situation and context. Therefore the meaning of social constructs varies over time and situation.

Such context-driven variation in interpretation has been demonstrated for all of the major ingredients of social impression. It has been shown that the same behaviors could have multiple interpretations depending on who performed them (e.g., Kunda & Sherman-Williams, 1993; Sagar & Schofield, 1980; Wojciszke, 1994), the interpretation of facial expressions could vary as a function of the situation (Trope, 1986), the interpretation of traits describing a person could vary as a function of other characteristics of that person (e.g., Asch, 1946; Asch & Zukier, 1984; Hamilton & Zanna, 1974; Kunda, Sinclair, & Griffin, 1995; Zanna & Hamilton, 1977), and the interpretation of stereotypes applied to a person can depend on other categories that the person belongs to (e.g., Kunda, Miller, & Claire, 1990) and on the person's traits (Deaux & Lewis, 1984).

It is important to note that in contrast to current serial accounts of impression formation (Brewer, 1988; Fiske & Neuberg, 1990), we do not give stereotypes a special processing role but rather treat them as no different from other information about people such as their traits and behaviors (cf. Smith & Zarate, 1992). Their influence on impressions, like that of any other information, depends on their patterns of association with other characteristics. For example, the impression of a Black person who excelled at college would be constrained both by the association between the stereotype of Blacks and academic ability (a negative association, in this case) and by the association between excelling at college and academic ability (a positive association). The ultimate impression of this person's ability will depend on the relative strengths of the direct positive and negative associations of the observed stereotyped category and behavior to ability as well as on their indirect associations to ability, by way of other nodes.

Model

Our theoretical claims do not specify how parallel constraint satisfaction occurs. Fortunately, it is now well understood how parallel constraint satisfaction can be performed by connectionist networks. We accordingly can refine our theoretical claims by offering the following computational model of impression formation.

Assumptions. Stereotypes, traits, and behaviors (characteristics) can be represented by units (nodes) in a connectionist network. Positive associations between two characteristics are represented by an excitatory link between the units that represent the two characteristics, and negative associations by an inhibitory link. For each characteristic that a particular person has been observed to have, an excitatory link is made between the unit representing that characteristic and a special, "observed" unit that is always active. In this manner, observed information is distinguished from inferred knowledge and is guaranteed to receive strong activation. This reflects the special status of observed information as information known to be true of the target person.²

The network goes through many repeated cycles of activation adjustment. In each, the activation of all units is adjusted in parallel, with the activation of each unit being updated on the basis of the activation of the units to which it is connected by excitatory and inhibitory links. Updating is repeated until all units have reached stable activation levels, that is, changes in their activation levels from one cycle to another are minimal. The network has then settled. After the network has settled, the activation level of a unit representing a characteristic reflects the extent to which the person is judged to have that characteristic. The overall impression formed of the person consists of the combined set of characteristics believed to characterize the person.

Similar models have been termed *interactive competition* models by McClelland and Rumelhart (1981) and *construction-integration* models by Kintsch (1988), but we class them generally as parallel-constraint-satisfaction models (see also Holyoak & Thagard, 1995, and Thagard, 1992). A connectionist model of this sort performs parallel constraint satisfaction by implementing all constraints as links and then settling the network to determine how the constraints can best be satisfied.

The process of impression formation. According to this model, impression formation proceeds as follows:

1. Observed information is activated. Not all the information available about the target person is observed and activated. Sometimes people may notice only salient group membership. In that sense, our model is compatible with those proposed by Brewer (1988) and by Fiske and Neuberg (1990), who argued that impressions are often based only on stereotypes. We would suggest, though, that it is also possible that sometimes people notice only behavior. Factors likely to influence which of the available pieces of information about a person is activated include the contextual salience of stereotype (S. Taylor, 1981) or behavior (Heider, 1958; Jones & Nisbett, 1972), the prior activation of constructs (Higgins & King, 1981), and perceiver goals (Brewer, 1988; Fiske & Neuberg, 1990).

2. Activation spreads from the activated observations to a fixed number of their immediate associates. This leads to an increase or a decrease in the activation level of these associates, depending on whether they have excitatory or inhibitory links with the original observations. The excitatory and inhibitory links and their strength come from the preexisting knowledge base. Here our model differs substantially from those proposed by Brewer (1988) and by Fiske and Neuberg (1990), who believe that in many cases stereotypes continue to dominate im-

pressions even after individuating information has been noticed. In contrast, we assume that, once activated, neither stereotypes nor individuating information have special status. Prior knowledge dictates the strength and direction of the connections between each type of construct and its associates. Both stereotypes and behaviors may vary in the strength of their associations with traits (Krueger & Rothbart, 1988), though we argue that behaviors are often associated with traits more strongly than are stereotypes.

3. The information is integrated by repeatedly updating the activation of all nodes until the network settles.

4. Additional inferences are made if necessary. The processes described so far reflect relatively automatic disambiguation of the meaning of incoming information. These processes are the focus of this article, and our program models only such automatic processes. Our theory also recognizes, however, that perceivers often engage in additional, more controlled inferential processes. Such additional causal inference is likely to occur if the perceiver cannot arrive at a coherent understanding of the information (Kintsch, 1988), if incoming information is particularly surprising (Kunda et al., 1990; Wong & Weiner, 1981), and if perceivers are motivated to obtain particularly thorough understanding of the target (Brewer, 1988; Fiske & Neuberg, 1990). The causal inferences may themselves be conducted through parallel-constraint-satisfaction processes. Thagard (1992) has shown how parallel constraint models may be applied to the evaluation of explanations, and Read and his colleagues have detailed how Thagard's model may be applied to explaining social behavior and to trait inference (Read & Marcus-Newhall, 1993; Read & Miller, 1993). Causal reasoning may also involve more serial hypothesis-testing processes (e.g., Trope, in press).

5. The products of inferences are integrated with previously activated knowledge to arrive at a final impression of the target person. This, too, is done through parallel constraint satisfaction. The degree to which the target person is assumed to have each characteristic is a function of each characteristic's ultimate level of activation.

The model just described is specific enough to be implemented in a running computer program that simulates how people form impressions of other people. The program simulates only the automatic aspects of impression formation. Details of this program, termed IMP (short for "IMPression formation") are provided in the Appendix.

Automatic and controlled processes in impression formation. Many important theoretical developments in social cognition over the last decade are predicated on the distinction between automatic processes, that is, relatively effortless processes carried out rapidly without conscious awareness or intention, and controlled processes, that is, more elaborate reasoning processes that require intention, awareness, and effort (e.g., Bargh, 1989; Devine, 1989; Gilbert, Pelham, & Krull, 1988). People's

² People may occasionally have trouble distinguishing observed from inferred knowledge, and may mistakenly believe that they have observed something which they have, in fact, only inferred (Johnson & Raye, 1981). We consider a characteristic to be observed if perceivers believe they have observed it.

impressions of other people often result from relatively automatic cognitive processes. It has been shown that trait inferences from behavior can occur spontaneously, without awareness or intent (Winter & Uleman, 1984; Winter, Uleman, & Cunniff, 1985). Two influential models of attribution (Gilbert, 1989; Trope, 1986, in press) assume that the identification of another person's behavior takes place automatically. For example, when a person is observed hitting another person, this behavior is automatically identified as aggressive. Gilbert and his colleagues (Gilbert, 1989; Gilbert et al., 1988) further assumed that the characterization of the person in trait terms (e.g., this person is aggressive) also takes place automatically. Various aspects of the context in which the behavior takes place can exert automatic influences on its identification and on the characterization of the person. For example, crying is automatically interpreted as reflecting sadness when it takes place at a funeral, but as reflecting joy when it takes place at a wedding (Trope, 1986, in press). Similarly, the activation of the stereotype of Blacks can automatically color the characterization of a person who had performed ambiguously aggressive behaviors (Devine, 1989).

Of course, the process of forming impressions of others is not always automatic. Often, people actively engage in more elaborate causal reasoning with the explicit intention of making sense of another's behavior. Current models of attribution assume that, following the initial automatic processes of impression formation, perceivers with sufficient cognitive resources may engage in more elaborate causal reasoning to assess whether the behavior was caused by the actor's traits or by alternative factors such as the constraints imposed by the situation (Gilbert, 1989; Trope, 1986, in press). In a similar vein, it has been suggested that people who do not themselves subscribe to culturally prevalent negative stereotypes of Blacks may be able to override the automatic impact of such stereotypes on their impressions of Black individuals if they have sufficient cognitive resources to engage in the required controlled processes (Devine, 1989).

In this article we propose that much of the cognitive work required to integrate stereotypes and individuating information is carried out automatically. Our program, IMP, simulates only such automatic processes. We show that many of the phenomena concerning how stereotypes and individuating information are integrated may be modeled by our program without invoking controlled inferences. We also point to a small number of phenomena illustrating that the social perceiver is not restricted to such automatic processes but can and does go beyond them to engage in controlled inferences under some circumstances. Because our aim is to demonstrate the scope of automatic parallel constraint processes in impression formation, we do not attempt to account for those phenomena that require more controlled causal reasoning. But we do wish to spell out the factors likely to trigger controlled processes. These include the following:

1. *The nature of the judgment task.* Some tasks explicitly require causal reasoning. These include the task of figuring out *why* a behavior took place, as in attribution studies that require participants to assess the extent to which different causal factors produced a behavior (e.g., Deaux & Emswiller, 1974). Another task that requires causal reasoning is the task of figuring out

whether an alleged behavior took place, as when a juror needs to decide whether a defendant is guilty or innocent (e.g., Bodenhausen, 1988). Task instructions that explicitly ask participants to identify causal relations among different pieces of information may also induce causal reasoning, even for tasks that could be handled automatically in the absence of such instructions (e.g., Asch & Zukier, 1984).

2. *The nature of information.* Perceivers may initiate causal reasoning when they have trouble understanding the information they have observed (Kintsch, 1988), or when they find it particularly surprising (Kunda et al., 1990; Wong & Weiner, 1981). This will occur if two or more of the observed pieces of information have strongly conflicting implications, as when target persons belong to seemingly conflicting social categories (e.g., Kunda et al., 1990) or behave in a manner that strongly violates their group's stereotype (e.g., Jussim, Coleman, & Lerch, 1987).

3. *Motivation.* Elaborate causal reasoning may be triggered when perceivers want to reach deep understanding of a given person or behavior (Brewer, 1988; Fiske & Neuberg, 1990).

We assume that phenomena involving judgment tasks that require causal reasoning are likely to result from controlled processes. The other conditions assumed to trigger controlled processes are often difficult to quantify; for example, we typically do not know just how surprised or confused participants were in the face of counterstereotypic information. In such cases, we assume that phenomena result from controlled processes if they contain other indicators of causal reasoning. One such indicator is the generation of novel constructs, an inferential feat unlikely to be accomplished automatically. Thus we assume that causal reasoning has taken place when perceivers creatively generate novel attributes that can explain how a person can be characterized by seemingly contradictory stereotypes (Hastie, Schroeder, & Weber, 1990; Kunda et al., 1990), or when perceivers construct new subtypes that can be used to isolate highly counterstereotypic individuals from their group's stereotypes (Kunda & Oleson, 1995; Weber & Crocker, 1983). Another indicator of causal reasoning is the presence of explicit causal language in protocols provided by participants (e.g., Asch & Zukier, 1984).

Knowledge representation. In our model, social knowledge is represented by the nature and strength of associations among constructs such as traits, behaviors, and social categories. For example, the belief that women are somewhat unassertive may be captured as a weak inhibitory link between *woman* and *assertive*, and the belief that straight As are highly indicative of academic ability may be captured by a strong excitatory link between these constructs. The simulations described later required that we make assumptions about knowledge representation, that is, about which other constructs each construct is connected to, and in what manner. The constructs that we represent are typically based on the independent and dependent measures and other stimulus information used in the research we are modeling. When we add constructs not directly represented in the materials and measures used in the research we are simulating, the added constructs are typically based on evidence from other research. For example, we assume that a well-dressed Black person may be viewed as more intelligent than a well-dressed White person because *well-dressed* invokes the subtype

of *Black businessmen*, which is viewed as particularly intelligent. Although the studies demonstrating such contrast effects (e.g., Jussim et al., 1987) do not themselves report measures of such subtypes, other research has documented their existence and their associations with the characteristics manipulated and assessed in the studies on contrast effects (Devine & Baker, 1991). In one case, to explain why individuating information that is not diagnostic of stereotypic traits can undermine the effects of stereotypes on impressions of targets (e.g., Locksley et al., 1982), we speculate that the nondiagnostic information invokes the stereotype of the ordinary, average person, which inhibits other stereotypes. This is the only case in which we hypothesize a construct whose existence and associations are based on our understanding of the domain rather than on empirical findings.

Connections among constructs are also inferred, in most cases, from empirical evidence. Associations between stereotypes and traits are typically based on findings suggesting that the trait is viewed as stereotypic (represented as an excitatory link) or counterstereotypic (represented as an inhibitory link) of the group in question. Similarly, associations between traits and behaviors are typically based on data indicating that the behavior is viewed as diagnostic (excitatory link), counterdiagnostic (inhibitory link), or nondiagnostic (no link) of the trait. In a small number of cases we inferred connections from our commonsense understanding of the domain in question, because the research we were simulating did not provide relevant empirical data. In all cases we inform the reader about the bases of our assumptions. The strength of associations is set at the default values unless otherwise specified. We deviate from the default values only when we have empirical evidence indicating that some connections are stronger than others. To readers concerned about the presence of some untested assumptions about the associations among constructs, we wish to point out that our model is not unique in making such assumptions; it is only unique in spelling them out. Other theories about how stereotypes are integrated with individuating information have not been presented at the level of detail that would make all their assumptions explicit.

We do not take a strong position about the hierarchy of constructs. We assume that stereotypes may be connected to behaviors directly (e.g., *housewife* inhibits *punches adults*) as well as indirectly, by way of traits (e.g., *Black* activates *aggressive*, which activates *violent shove*). The knowledge representations that we use in our simulations are meant to be viewed as schematic, that is, as illustrations that it is possible to simulate empirical results under a reasonable set of assumptions about knowledge representation. Therefore our simulations incorporate only the knowledge that seems central to the phenomenon under consideration; we do not attempt to include all the associates of the constructs in question. Typically, various reasonable modifications of these knowledge structures (e.g., adding or deleting an intermediate trait node between a stereotype and a behavior) would yield comparable patterns of results.

We now review the literature on the effects of stereotypes and individuating information on impressions of individuals. We describe the key phenomena to emerge from this literature and show how our model can account for them and how IMP can

simulate them. We begin by discussing how stereotypes and individuating information constrain each other's meaning. We then turn to a discussion of how these sources of information are combined to form impressions of individuals.

Stereotypes and Individuating Information Constrain Each Other's Meaning

The first three phenomena concern how stereotypes affect the interpretation of behavior and traits and how individuating information affects the interpretation of stereotypes.

Phenomenon 1: Stereotypes color the meaning of behavior. When a person known to belong to a particular social category performs an ambiguous behavior, the stereotype associated with that category can color the meaning of the behavior. Thus, a shove was viewed as more violent when performed by a Black than by a White (Duncan, 1976; Sagar & Schofield, 1980), a mixed performance on an academic test was considered better when performed by a child coming from a high versus low socioeconomic background (Darley & Gross, 1983; Kameda, 1985), and the behavior *hit someone who annoyed him or her* was interpreted as *punched an adult* when performed by a construction worker but as *spanked a child* when performed by a housewife (Kunda & Sherman-Williams, 1993).

In our model, such stereotype-driven construal of behavior can occur because the stereotype activates a trait which, in turn, activates some of the possible interpretations of the behavior, or because the stereotype activates some interpretations of the behavior directly. The indirect route is depicted in Figure 1. Recall that each piece of information that has been observed, that is, that the perceiver has seen or been told about, is connected to a special unit named *observed*. The remaining connections are all based on the perceiver's preexisting beliefs. Each of the units depicted also has many additional associates that are not portrayed in the figures or used in our simulations.

As discussed earlier, when a person observed pushing someone is Black, *Black* activates *aggressive*, which, in turn, activates *violent push*, whose final activation in our IMP simulation is .57. In contrast, when the person is White, *aggressive* does not get activation from the stereotype node, and so *violent push* ends up with less activation (-.45). In this manner the behavior *pushed someone* can come to be interpreted differently as a function of the stereotype associated with the person who performed it (Sagar & Schofield, 1980).

Stereotypes could also affect the meaning of behavior through direct associations with relevant behaviors. Such effects could be indicated by adding an excitatory connection between *Black* and *violent push*, and, possibly, an inhibitory connection between *White* and *violent push*. This would strengthen the effects of stereotypes on the construal of behavior. We expect that in some cases the effects of stereotypes on the meaning of individuating information arise from direct links between the stereotype and associates of that information, in some cases from indirect links through stereotypic attributes associated with that information, and in some cases through a combination of direct and indirect associations.

Phenomenon 2: Stereotypes color the meaning of traits. Our model led to the prediction that when a trait is used to describe




Figure 2. Stereotypes influence the meaning of the trait *aggressive* (Phenomenon 2). The network on the left activates *argue* (.39) more than *punch* (-.09), whereas the network on the right activates *punch* (.53) more than *argue* (.28).

a member of a stereotyped group, its meaning can be influenced by the group's stereotype. The same trait can imply different behaviors when applied to members of different groups. In a series of studies generated to test this prediction, Kunda, Sinclair, and Griffin (1995) found that perceivers rated lawyers and construction workers as about equally aggressive but nevertheless held very different expectations about their likely aggressiveness-related behaviors: Lawyers were viewed as more likely to argue and complain and as less likely to punch and yell insults than were construction workers.

In our model, such stereotype-driven construal of traits can occur because other traits associated with the stereotype activate some of the associates of the trait in question. Kunda et al. (1995) showed that, in addition to being aggressive, lawyers are also viewed as upper middle class and verbal, whereas construction workers are viewed as working class and unrefined. These additional aspects of the stereotype are themselves associated with aggressive behaviors and so can constrain the meaning of aggressiveness when this trait is applied to lawyers or construction workers. Such constraints are depicted in Figure 2. It may be seen that *lawyer* and *construction worker* activate *aggressive* to the same extent. But *lawyer* also activates *upper middle class* and *verbal*, whereas *construction worker* also activates *working class* and *unrefined*. These additional associates of the stereotypes have different patterns of association with the behavioral associates of *aggressive*: *Upper middle class* inhibits *punch*, whereas *working class* activates it. *Verbal* activates *argue*, whereas *unrefined* does not. The end result is that in the case of the lawyer, *punch* gets less activation and *argue* gets more activation than in the case of the construction worker.

Phenomenon 3: Individuating information can determine which of the stereotype's subtypes is used. When a stereotyped person has counterstereotypic characteristics, these may constrain the meaning of the stereotype. Many groups are stereotyped as having distinct subtypes. For example, the group *elderly women* is believed to include the subtype of *grandmother*, viewed as kindly and cheerful, and the subtype of *senior citizen*, viewed as lonely and worried (Brewer, Dull, & Lui, 1981). Similarly, the group *Blacks* is believed to include the subtype of *ghetto Blacks*, viewed as hostile, poor, jobless, and negative in personality, and the subtype of *Black businessmen*, viewed as ambitious, intelligent, upwardly mobile, and positive in personality (Devine & Baker, 1991). The individuating information

that characterizes a member of such a multifaceted group may determine which of the subtypes of the stereotypes will be activated. For example, a well-dressed, well-spoken Black man may be thought of as a Black businessman rather than as a ghetto Black (cf. Brewer, 1988; Fiske & Neuberg, 1990). Accordingly, Deaux and Lewis (1984) found that women engaging in male role behaviors (e.g., head of household) and men engaging in female role behaviors (e.g., manages the house) were viewed as substantially more likely to be gay than were men and women engaging in stereotypic role behaviors; apparently the counterstereotypic behaviors activated the gay subtypes of men and women. Likewise, Costrich, Feinstein, Kidder, Marecek, and Pascale (1975) found that men and women behaving in counterstereotypic manner (passively for men, aggressively for women) were considered to be in greater need of psychotherapy than were those engaging in more stereotypic behavior; the counterstereotypic behaviors appear to have activated the abnormal subtypes of men and women. In a similar vein, Fiske, Neuberg, Beattie, and Milberg (1987) reported that subjects' open-ended descriptions revealed that some subjects categorized individuals with counterstereotypic attributes (e.g., a rowdy professor) as belonging to distinct subtypes of their group (e.g., the charismatic lecturer type). Studies focusing on the effects of counterstereotypic individuals on stereotypes have also found that some counterstereotypic individuals appear to be categorized as belonging to atypical subtypes of their groups (Rothbart & Lewis, 1988; Weber & Crocker, 1983).

Our model accounts for the constraining effects of individuating information on the meaning of stereotypes by assuming that the individuating information activates some of the stereotype's subtypes and inhibits others. Consider the stereotype of Blacks, depicted in Figure 3. We assume that the subtype *ghetto Black* is more strongly associated with *Black* than is the subtype *Black businessman*. This assumption is based on data showing that there is considerable overlap between the global stereotype of Blacks and the subtype *ghetto Blacks* but practically no overlap between the global stereotype and the subtype *Black busi-*

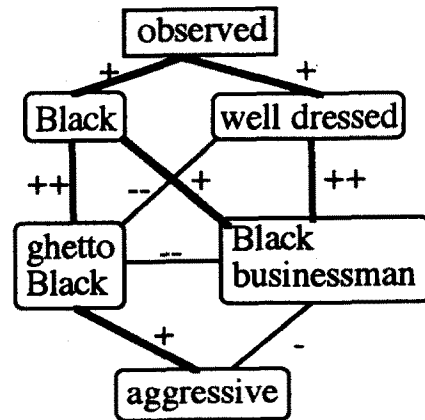


Figure 3. Individuating information affects activation of a stereotype's subtypes (Phenomenon 3). This network deactivates *ghetto Black* and activates *Black businessman*, but if *well dressed* is not connected to *observed* the reverse is true.

Table 1
Activation Level of Subtypes of Black and of the Trait Aggressive in the Presence and the Absence of the Individuating Information Well Dressed About a Black Person

Individuating information	Ghetto Black	Black businessman	Aggressive
None	.83	-.77	.60
Well dressed	-.79	.79	-.61

nessmen (Devine & Baker, 1991). In the absence of individuating information, *ghetto Black* will therefore receive stronger activation than *Black businessman*, as seen in Table 1. However, the individuating information *well dressed* inhibits *ghetto Black* while activating *Black businessman*. Therefore, in the presence of this information, *Black businessman* will be more strongly activated than *ghetto Black*.

The rightmost column of Table 1 depicts the activation levels of a trait stereotypic of Blacks, *aggressive*. It is apparent that this trait is not nearly as strongly activated in the presence of the information that the person is well dressed than in its absence, because *aggressive* is inhibited by the subtype *Black businessman*, which is activated by the individuating information. This may be one route through which individuating information may undermine the effects of stereotypes on trait ratings (see Phenomenon 5): Individuating information, both diagnostic and nondiagnostic of a given stereotypic trait, may undermine the effects of the global stereotype on that trait by activating an atypical subtype of the stereotyped group.

This analysis shows how individuals who violate their groups' stereotypes can be classified as belonging to preexisting atypical subtypes of these groups (Hewstone, 1994; Rothbart & Lewis, 1988; Weber & Crocker, 1983). Sometimes exposure to counterstereotypic individuals leads perceivers to construct novel subtypes to which these individuals can be assigned (Kunda & Oleson, 1995). To account for such constructive phenomena it is necessary to call on controlled processes.

Combining Multiple Sources of Information to Form Impressions of Individuals

So far we have reviewed findings indicating that the meaning of a person's behavior and traits can be affected by the stereotypes of groups to which that person belongs, and that the meaning of such stereotypes can be affected by the person's behavior. We showed how our model of parallel constraint satisfaction can account for these processes. We next turn to examining how stereotypes and individuating information may be combined to form impressions of individuals. We show how the processes of mutual constraints on meaning that we have described can be applied to understanding the key phenomena found in this literature.

Stereotypes Without Individuating Information

Phenomenon 4: Stereotypes in the absence of individuating information color impressions. When all one knows about an

individual is that he or she belongs to a particular social category, the stereotype associated with that category colors one's impression of the individual. Thus, a person described only by a male name was viewed as more assertive than a person described only by a female name (Locksley, Borgida, Brekke, & Hepburn, 1980; Rasinski, Crocker, & Hastie, 1985) and as more likely to be characterized by a host of traits associated with masculinity (Heilman, 1984). Similarly, a person described only as a night person was viewed as more unpredictable than a person described only as a day person (Locksley et al., 1982). Students described only as having different majors were expected to behave differently in everyday situations as well as in psychology experiments (Nisbett, Zukier, & Lemley, 1981). Finally, Blacks portrayed only in photographs were viewed as more superstitious, lazy, emotional, untidy, and immoral than were Whites portrayed in photographs (Secord, Bevan, & Katz, 1956), and a child viewed on a brief videotape in surroundings implying high socioeconomic background was viewed as possessing better liberal arts skills, cognitive ability, and emotional maturity than she was when viewed in surroundings implying low socioeconomic class (Darley & Gross, 1983). People may sometimes be reluctant to apply a stereotype to an individual about whom they know nothing else, perhaps because they feel it is inappropriate to do so (Darley & Gross, 1983). But there is ample evidence that stereotypes can and do color one's beliefs about the likely traits and behaviors of individuals about whom no additional individuating information is available.

Our model accounts for this in a very straightforward manner. When a stereotype is activated, activation spreads to the traits associated with it. For example, a construction worker might be expected to be aggressive whereas a housewife might be expected to be unaggressive. Activation also spreads to the behaviors associated with these traits, directly from the stereotype as well as indirectly, through the activated traits. Thus stereotype-related traits and behaviors get strong activation and so are viewed as likely for the person characterized by the stereotype.

Stereotypes With Diagnostic Individuating Information

The effects of stereotypes in the presence of diagnostic individuating information, that is, information relevant to the judgment in question, depend on the kind of judgment being made, the ambiguity of the individuating information, and the certainty with which it is believed to characterize the target. We first examine the effects of stereotypes in the presence of individuating information assumed to be certain and then examine their effects in the presence of less certain information, of the sort encountered by juries. We begin by examining the effects on trait ratings and then show how other judgment tasks, such as behavioral predictions, can yield different patterns of results.

Phenomenon 5: Stereotypes typically do not affect trait ratings in the presence of unambiguous diagnostic information. Even though stereotypes affect trait ratings of individuals about whom nothing else is known, stereotypes typically have no such effects when the individual is also known to have engaged in an unambiguous behavior that is clearly diagnostic of the trait in question. Thus a person identified only as Tom was expected to

be more assertive than a person identified as Nancy. However, when each was known to have abruptly interrupted another student so as to break into a discussion, a clearly assertive behavior, they were expected to be equally assertive (Locksley et al., 1980). This pattern has been widely replicated in studies that use traits as the dependent measure and assess the impact of unambiguous behavior that is assumed to have occurred. Thus, one or more clear and diagnostic behaviors by a target person have been shown to undermine the effects of the stereotypes of night and day people on ratings of the target's predictability (Locksley et al., 1980), the effects of the stereotypes of construction worker and housewife on the target's aggressiveness (Krueger & Rothbart, 1988, Experiment 2; Kunda & Sherman-Williams, 1993), the effects of gender stereotypes on the target's gender-related traits (Berndt & Heller, 1986; Deaux & Lewis, 1984; Dipboye & Wiley, 1977; Glick, Zion, & Nelson, 1988; Heilman, 1984; Rasinski et al., 1985), the effects of racial stereotypes on the target's competence (Jussim et al., 1987, for participants exposed to lower class targets; Linville & Jones, 1980, for participants exposed to targets with weak credentials), the effects of stereotypes of social class on the target's academic ability (Kameda, 1985, for participants exposed to high- and low-ability targets), the effects of stereotypes of skinheads and hairdressers on the target's aggressiveness (Macrae, Sheperd, & Milne, 1992, when information was provided by a high-credibility source), the effects of stereotypes of car salesmen and librarians on the target's extraversion (Kunda & Sherman-Williams, 1993), and the effects of the stereotypes of homosexuals and heterosexuals on the target's creativity (Jussim, Nelson, Manis, & Soffin, 1995). It is clear from this set of findings that the target's behavior has been shown to undermine the effects of stereotypes based on all the major social categories—race, gender, class, profession, and personality type—on the target's trait ratings.

Our model accounts for this pattern by assuming that the association between behaviors and traits is stronger than the association between stereotypes and traits. Therefore, in the presence of a behavior that is strongly associated with a trait, the effects of stereotypes that are less strongly associated with that trait can be overwhelmed. We have strong empirical grounds for making this assumption. Many studies on the effects of stereotypes have also varied the nature of the target's behavior. Almost invariably, such variations in individuating information have massive effects on trait ratings. Thus variations in behaviors reflecting aggressiveness, assertiveness, extraversion, masculinity or femininity, academic success, professional accomplishment, and social class have all been shown to have massive impact on how the targets were rated on the traits that these behaviors imply (Biernat & Manis, 1994; Deaux & Taynor, 1973; Dipboye & Wiley, 1977; Feldman-Summers & Kiesler, 1974; Glick et al., 1988; Heilman & Stopeck, 1985; Jackson, Sullivan, & Hodge, 1993; Jussim et al., 1987; Kameda, 1985; Krueger & Rothbart, 1988; Kunda & Sherman-Williams, 1993; Linville & Jones, 1980; Locksley et al., 1980; Macrae et al., 1992, when information was provided by a high-credibility source; Pratto & Bargh, 1991; Triandis, Loh, & Levin, 1966). Similarly, variations in targets' behavioral roles and personality traits have been shown to have substantial impact on how the targets were rated on re-




Figure 4. Stereotypes do not affect trait ratings in the presence of unambiguous diagnostic behavior (Phenomenon 5). See Table 2 for values of *aggressive*, which are similar for both networks.

lated traits (Deaux & Lewis, 1984; Del Boca & Ashmore, 1980; Grant & Holmes, 1981). Individuating information had substantial impact on trait ratings in every study in which it has been varied.

A meta-analysis revealed that the average effect size for individuating information in these studies ($r = .71$) was substantially larger than the average effect size typically obtained for stereotypes even in the absence of individuating information ($r = .25$), as revealed in a meta-analysis conducted on the studies cited in the discussion of Phenomenon 4.³

Figure 4 portrays a network that incorporates the assumption that behavior activates traits more strongly than stereotypes do. It depicts a perceiver who has observed a member of a stereotyped group (housewife or construction worker) perform an unambiguously aggressive behavior (punched a neighbor). Table 2 depicts the activation levels of *aggressive* when the stereotype is observed in the presence and in the absence of that behavior. It is apparent that the results successfully simulate the empirical findings: The stereotypes exert substantial impact on *aggressive* in the absence of the behavior (i.e., there is a large difference in the activation of *aggressive* when the target is a construction worker vs. a housewife), but the stereotypes exert practically no

³ We located articles for this meta-analysis and the one reported later in this article by checking all the relevant references cited in the literature reviews cited in this article (which include several recent and comprehensive reviews of the field: Brewer, 1988; Hamilton & Sherman, 1994; Fiske & Neuberg, 1990) as well as all the relevant references cited in all the empirical articles that we cite. This was supplemented by a computerized search through the PsychLit database, 1974–1995, using the key word *stereotype*. We included only studies reported in published journal articles, in which adult participants evaluated particular target persons (rather than average or typical group members or targets that participants were asked to imagine). Studies had to use at least two targets who belonged to differently stereotyped groups (e.g., a male and a female target) for which participants held preexisting stereotypes (we excluded studies using experimentally induced stereotypes of previously unknown groups). Effect sizes (r) were calculated from statistics provided in the cited articles, using the conversion formulas provided by Rosenthal (1991). Effects described only as nonsignificant were assigned values of 0.00 for effect size (Rosenthal, 1991). When a study reported more than one measure of impressions of target, these measures were averaged so that each study yielded a single effect size. To obtain the overall average effect size, r s were converted into Fisher's z 's, weighted by the sample sizes on which they were based, averaged, and reconverted into r s.

Table 2
Activation Level of Aggressive for Different Stereotypes in the Presence of No Individuating Information, an Unambiguous Aggressive Behavior, and an Ambiguous Behavior

Stereotype	No individuating information (Phenomenon 4)	Unambiguous behavior: punched adult (Phenomenon 5)	Ambiguous behavior: hit someone (Phenomenon 6)
Housewife	-.38	.69	-.69
Construction worker	.27	.68	.73
Difference	.65	-.01	1.42

impact on *aggressive* in the presence of the behavior. As seen in Figure 4, this simulation includes direct links between the stereotypes (*construction worker* or *housewife*) and the behavior (*punch adult*). Similar results were obtained when these links were eliminated. In IMP's simulations, the difference in aggression between construction worker and housewife gets smaller as the weight on the link between *punch adult* and *aggressive* is increased, that is, as the behavior becomes more diagnostic of the trait.

Note that we assume that behaviors are associated more strongly than stereotypes with impressions because that has been true for the stereotypes and behaviors investigated to date, on average. We also assume, though, that both behaviors and stereotypes can vary considerably in the strength of their associations with different traits. When particularly strong stereotypes are combined with particularly weak behaviors, we would not expect the effects of stereotypes on impressions to be undermined by the behaviors (Beckett & Park, 1995; Krueger & Rothbart, 1988).

Phenomenon 6: Stereotypes do affect trait ratings in the presence of ambiguous individuating information. When members of stereotyped groups perform behaviors that are diagnostic of a given trait but that are ambiguous and open to multiple construals, stereotypes do affect the ratings of the target person on that trait. Thus a construction worker who performed an ambiguously aggressive behavior was rated as more aggressive than a housewife who performed that behavior (Krueger & Rothbart, 1988; Kunda & Sherman-Williams, 1993), a child from a high socioeconomic background who displayed a mixed performance on an ability test was rated as higher in ability than a child from a low socioeconomic background with the same performance (Darley & Gross, 1983; Kameda, 1985), and a mental patient describing himself in an ambiguous manner was viewed by analytical therapists as more disturbed than a job applicant describing himself in the same manner (Langer & Abelson, 1974). Whenever traits, which are notoriously open to multiple construals (Asch, 1946; Asch & Zukier, 1984; Griffin & Ross, 1991; Hamilton & Zanna, 1974; Zanna & Hamilton, 1977), have been used as individuating information, stereotypes affected the ratings of targets on related traits (Deaux & Lewis, 1984; Del Boca & Ashmore, 1980; Grant & Holmes, 1981).

It has been suggested that when stereotypes affect impressions in the presence of ambiguous individuating information, they do so by affecting the construal of that information (Kunda & Sherman-Williams, 1993). Thus a construction

worker known to have hit someone who annoyed him is viewed as more aggressive than a housewife known to have hit someone who annoyed her because the behavior is interpreted differently for the two. The construction worker is assumed to have punched an adult, whereas the housewife is assumed to have spanked a child. Such stereotype-driven construals of behavior have been shown to mediate the effects of stereotypes on ratings of targets (Kunda & Sherman-Williams, 1993).

Figure 5 depicts a network in which the perceiver observed a stereotyped target engage in an ambiguous behavior. In this figure the stereotype is connected directly to each of the different behavioral associates of *hit someone* and so constrains the meaning of that behavior. Similar results were obtained when the stereotype was connected to these behavioral associates only indirectly, through additional aspects of the stereotype: *housewife* activated *takes care of children*, which activated *spank child*; and *construction worker* activated *rough work place*, which activated *punch adult*. The rightmost column of Table 2 depicts the activation levels of *aggressive* when the target's behavior was ambiguous. It is apparent that our model successfully simulated experimental findings: Although the stereotype did not affect the activation of *aggressive* when the behavior was unambiguous, it did when the behavior was ambiguous.

Phenomenon 7: Stereotypes can provoke contrast effects on trait ratings. When a member of a stereotyped group is known to have engaged in a clear, diagnostic, counterstereotypic behavior, stereotypes typically have no effects on trait ratings, as discussed earlier. However, a small number of studies found what has been termed a *contrast effect*: Individuating information suggesting a counterstereotypic trait leads targets to be viewed as particularly extreme on that trait. Thus, although




Figure 5. Stereotypes affect trait ratings in presence of ambiguous diagnostic behavior (Phenomenon 6). The construction worker is judged to be much more aggressive than the housewife (see Table 2).

Blacks are typically viewed as less academically competent than Whites, Blacks with strong academic credentials are often viewed as even more competent than Whites with comparable credentials (Jackson et al., 1993; Jussim, et al., 1987; Linville & Jones, 1980). Likewise, for dimensions on which men are typically viewed as more competent than women, a man with low ability may be judged as even less competent than a woman with comparable ability (Deaux & Taynor, 1973), and a woman displaying masculine behaviors may be viewed as even more manly than a man displaying the same behaviors (Costrich et al., 1975; Pratto & Bargh, 1991). In a similar vein, although paraplegics are typically viewed more positively than are schizophrenics, a schizophrenic characterized by a mixture of positive and negative information may be viewed as more likable than a similarly characterized paraplegic (Fiske & Von Hendy, 1992, for participants who were low self-monitors).

Our model can account for these findings by assuming that when a member of a stereotyped group (Blacks) has a counterstereotypic attribute (well dressed), this leads to the activation of a preexisting subtype (Black businessman) that is viewed as unusually high on the dimension in question (intelligence). This example is based on research documenting the existence of a distinct subtype of Black businessmen (Devine & Baker, 1991). Whereas Blacks were described by the majority of participants in that study as unintelligent and were not described as well dressed by any participant, Black businessmen were described by the majority as intelligent and were described by almost half of the participants as well dressed. Figure 3 portrays how the counterstereotypic attribute *well dressed* may activate the atypical subtype of Blacks, *Black businessmen*. IMP can simulate a contrast effect in which the well-dressed Black person will be viewed as more intelligent than a similarly well-dressed White person if one further assumes that the subtype of *Black businessmen* is more strongly associated with intelligence than is *White* or any of the subtypes of Whites activated by *well dressed*.

In this manner our model can readily account for many contrast effects. However, it seems unlikely that extreme subtypes exist in all cases where contrast effects occur. Therefore we also assume that contrast effects can occur even in the absence of preexisting extreme subtypes, through causal reasoning. One may reason, for example, that to achieve strong academic records, Black people need to overcome many obstacles and so need to be more competent than White people achieving the same records.

There are too few demonstrations of contrast effects to permit an assessment of the question of when stereotypes will provoke contrast effects rather than the more frequently found null effects on trait ratings. If we are correct in assuming that contrast effects often result from causal reasoning, it seems likely that such effects may occur when the individuating information is extremely counterstereotypic. Such extremity is likely to trigger surprise, which, in turn, will lead to controlled processes.

Phenomenon 8: The effects of stereotypes in the presence of diagnostic individuating information depend on the judgment task. As we have seen, when subjects are asked to make trait ratings about a stereotyped individual known to have engaged in an unambiguous trait-related behavior, the stereotype typi-

cally has practically no impact on these ratings. A different pattern emerges when subjects given comparable information are asked to predict the target's future trait-related behavior or to assess the target's suitability to jobs requiring trait-related skills. Here, stereotypes do affect judgment in most cases: Members of groups stereotyped as more likely to have a given trait are viewed as more likely to engage in trait-related behaviors and as more suitable for trait-related jobs even in the presence of clear individuating information. Thus ethnic stereotypes influenced judgments about whether a target known to have engaged in a crime would repeat the crime (Bodenhausen & Wyer, 1985, Experiments 1 & 2), racial stereotypes affected predictions about how well a target with strong academic credentials would do in college (Jackson et al., 1993, Experiments 1 & 2, for participants exposed to a high-ability target), gender stereotypes affected judgments about the suitability for male and female jobs of targets with clear academic records (Dipboye, Fromkin, & Wiback, 1975, for participants exposed to a high-ability target; Heneman, 1977) and with clear work background and experience (Glick et al., 1988; Rosen & Jerdee, 1974a, 1974b), and gender stereotypes affected predictions about the behavioral choices of targets who had previously chosen to engage in stereotype-consistent or inconsistent activities (Berndt & Heller, 1986, for college age participants) as well as predictions about whether targets would behave assertively or passively in psychological experiments (Beckett & Park, 1995, targets with photos). Although there are several studies in which stereotypes did not affect behavioral predictions in the presence of individuating information (Beckett & Park, 1995, for participants exposed to targets without photos; Heneman, 1977, for participants exposed to low- and medium-ability targets; Jackson et al., 1993, for participants exposed to a low-ability target; Locksley et al., 1982), assimilation of the targets to the stereotypes is considerably more common for studies using behavioral predictions as the dependent measure than for studies using traits as the dependent measure.

Different judgment tasks have been shown to yield different patterns of stereotypic effects even when applied to the same target within a single study (Jackson et al., 1993). Several studies have found that stereotypes affect behavioral prediction about a target even when they have no impact on the trait ratings of the same target (Berndt & Heller, 1986, for college-age subjects; Deaux & Lewis, 1984, for ratings of feminine traits; Glick et al., 1988).

In a series of studies inspired by the current model, Kunda et al. (1995) showed that when differently stereotyped targets performed a behavior diagnostic of a given trait, this behavior undermined the effects of the stereotypes on ratings of the targets on that trait. Nevertheless, stereotypes still affected predictions about the targets' likelihood of engaging in other behaviors related to that trait. Thus a construction worker and an accountant were judged as equally unaggressive after performing a clearly unaggressive behavior, failing to react to an insult. Nevertheless, the construction worker was still viewed as more likely to engage in aggressive behaviors such as punching and cursing than was the accountant.

We discussed in an earlier section how differently stereotyped targets who are rated the same on a given trait can nevertheless




Figure 6. Stereotypes can affect behavioral predictions even when they do not affect trait ratings (Phenomenon 8). See Table 3 for final activations of *aggressive* and *punch*. In the no-individuating-information conditions, the *walked away* unit is omitted.

be expected to display different behavioral manifestations of that trait (Phenomenon 2). The same logic applies when a stereotypic trait has been undermined by individuating information: When targets' trait-related behaviors undermine the effects of stereotypes on how the targets are rated on that trait, stereotypes may nevertheless affect predictions about which other trait-related behaviors the targets are likely to engage in. This will occur if these behaviors are also associated with additional stereotypic traits that have not been undermined by the targets' behavior and so can continue to exert influence on the behavioral predictions. Accordingly, Kunda et al. (1995) found that the continuing impact of stereotypes on behavioral predictions even after their impact on traits had been undermined occurred because the behavioral predictions were also influenced by another stereotypic dimension, social class, which had not been undermined by the individuating information. The now unaggressive but still working-class construction worker was viewed as more likely to punch another than was the equally unaggressive but upper-middle-class accountant, because such punching was considered more likely for working-class than for upper-middle-class people. Indeed, the effect of stereotypes on the prediction of behaviors was greater for those behaviors more strongly associated with social class.

Figure 6 portrays a network in which stereotypes affect trait-related predictions about a target's behavior even though they do not affect how the target is rated on that trait. It may be seen that *accountant* has an inhibitory link with *aggressive* whereas *construction worker* has an excitatory link with *aggressive*. *Aggressive*, in turn, has an excitatory link with *punch*. Therefore, as seen in Table 3, in the absence of individuating information, both *aggressive* and *punch* are more strongly activated for the construction worker than for the accountant. When the target is observed performing an unaggressive behavior, failure to react to an insult, this behavior strongly inhibits *aggressive*. As seen in Table 3, *aggressive* receives comparable low activation for the construction worker and for the accountant—the effect of the stereotype on ratings of this trait has been undermined. But the stereotype nevertheless continues to affect the activation level of the aggressiveness-related behavior *punch* because *punch* is also associated with other aspects of the stereotypes that are relatively unaffected by the behavior.

In this manner our model can readily account for the finding that, in the presence of diagnostic individuating information,

stereotypes can affect predictions about the target's probable trait-related behavior even though they do not affect how the target is rated on that trait. In contrast, this finding poses a serious problem for the serial models proposed by Brewer (1988) and by Fiske and Neuberg (1990) because those models provide no reason for distinguishing among different judgment tasks.

Another judgment task that may be affected by stereotypes is causal attribution. Here, subjects are asked to decide what caused the target's behavior, typically by rating the extent to which the behavior was caused by the target's internal dispositions and by external, situational factors. Stereotypes typically do affect attributions, such that behaviors are more likely to be viewed as caused by internal dispositions when consistent with a stereotype than when inconsistent with it. For example, weak academic credentials were attributed more to ability when characterizing a Black person than a White person (Jackson et al., 1993), and successful performance on a masculine task was attributed more to ability for a male than for a female (Deaux & Emswiller, 1974). (For similar results, see also Duncan, 1976; Feather & Simon, 1975; Hewstone & Ward, 1985; Macrae & Shepherd, 1989; D. Taylor & Jaggi, 1974; Yarkin, Town, & Wallston, 1982). Analysis of the relative impact of the different potential causes on behavior is a task that, by its very nature, demands controlled, inferential processes (Trope, 1986, in press). In so arguing, we do not wish to challenge the view that the early stages of attribution can be carried out automatically (Gilbert, 1989). The judgments performed in these early stages are the identification of behavior and the characterization of actors in trait terms, tasks that we, too, assume can be automatic. But the attribution of behaviors to different causes is a different kind of task, more laden with causal reasoning than are trait characterizations. Indeed, subjects appear to make attribution judgments more slowly than they make trait judgments, suggesting that attribution judgments are less automatic than trait ratings (Smith & Miller, 1983). Therefore, the effects of stereotypes on causal attribution cannot be accounted for by the more automatic processes simulated by IMP and are better accounted for by higher order causal reasoning of the sort modeled by Read and his colleagues (Read & Miller, 1993) or of the sort proposed by Trope (1986, in press).

Phenomenon 9: Stereotypes sometimes affect impressions in the presence of uncertain information. So far, we have dis-

Table 3
Activation Level of *Aggressive* and *Punch* for Different Stereotypes in the Absence of Individuating Information and in the Presence of the Diagnostic Individuating Information Walked Away From an Insult, an Unaggressive Behavior (Phenomenon 8)

Stereotype	No individuating information		Unaggressive behavior	
	Aggressive	Punch	Aggressive	Punch
Accountant	-.58	-.71	-.80	-.72
Construction worker	.50	.60	-.70	-.04
Difference	1.08	1.31	.10	.68

cussed the effects of stereotypes on impressions in the presence of diagnostic information that is assumed to be true—behaviors that the target is known to have performed, or traits and background information known to characterize the target. We next turn to a different kind of judgmental task, one in which perceivers need to determine whether a member of a stereotyped group had performed a given behavior, in the presence of inconclusive evidence. This is the task facing a juror trying to determine whether or not a defendant is guilty. Many other everyday judgments are likely to be based on uncertain evidence, but, with one exception (Macrae et al., 1992, for participants exposed to a low-credibility source), all of the research in which participants were given uncertain evidence about whether or not a behavior occurred examined assessments of whether or not the target was guilty of criminal behavior. Research on the effects of stereotypes on such assessment of guilt is typically discussed in the same breath as research on the effects of stereotypes on impressions in the presence of certain individuating information, and findings from research on guilt assessment are assumed to have implications for the question of how stereotypes are combined with individuating information assumed to be certain. But from our perspective, the two kinds of judgment tasks are inherently different. Whereas stereotypes and certain individuating information may be integrated in a relatively automatic way, through spreading activation, the task of determining whether or not a behavior occurred seems to require reliance on more controlled processes involving causal reasoning and explanation. A juror needs to evaluate the credibility of each piece of evidence and decide whether the full set is best explained by innocence or by guilt. Models of such explanatory reasoning have been developed by Thagard (1989) and by Read and Miller (1993). IMP models only a much simpler kind of automatic inference pattern in which the excitatory and inhibitory associations of a stereotype produce judgments.

Research examining the effects of stereotypes on guilt judgments made by subjects operating under normal conditions (i.e., no unusual mood or stress; we discuss the effects of these factors on judgment in a later section) has yielded mixed results. Several studies have found that stereotypes do affect judgments, so that members of groups stereotyped as likely to perform a given crime are viewed as more likely than members of other groups to be guilty of that crime (Bodenhausen, 1988; Bodenhausen & Lichtenstein, 1987; Klein & Creech, 1982; Ugwuegbu, 1979). However, there are also several studies involving comparable crimes in which comparable stereotypes had no effects on judgments of guilt likelihood (Bodenhausen, Kramer, & Susser, 1994; Bodenhausen, Sheppard, & Kramer, 1994; Ugwuegbu, 1979, for participants given strong and weak evidence).

It is difficult to tell from existing research when stereotypes will affect guilt judgments and when they will not. There is some evidence that stereotype effects may be more likely when the evidence is ambiguous, rather than strong or weak (Ugwuegbu, 1979). However, most studies that obtained no stereotype effects used ambiguous evidence, consisting of items implying guilt as well as items implying innocence. It remains possible that slight increases in the extent to which the evidence appears conclusive may suffice to eliminate stereotype effects. It has also

been suggested that stereotype effects may be more likely when subjects process the evidence with the goal of performing a complex task such as guilt assessment than when they process it with the goal of performing a simpler task such as trait evaluation (Bodenhausen & Lichtenstein, 1987). But several studies in which subjects expected to make guilt judgments also failed to yield stereotype effects (Bodenhausen et al., 1994). Some of the effects of stereotypes on guilt judgments may be due to the effects of the stereotypes on interpreting the evidence, as suggested by the finding that the effects of stereotypes were eliminated when the stereotype was presented only after the evidence and so could not affect its interpretation (Bodenhausen, 1988).

In sum, further research is necessary to determine when stereotypes affect judgments of guilt. It is important to note, however, that this kind of judgment is quite different from the judgment that subjects are called on to make as they form impressions of people known to have performed certain behaviors in that it is much more likely to require controlled processes. Therefore, the findings about how stereotypes affect guilt judgments may have little implication for the question of how stereotypes are integrated with individuating information that is viewed as certain.

Stereotypes and Nondiagnostic Individuating Information

So far, we have discussed the effects of individuating information that is diagnostic of the judgments that subjects are asked to make. There has also been considerable interest in how stereotypes affect judgments of targets in the presence of individuating information that is nondiagnostic of these judgments, that is, irrelevant to them. Following Hilton and Fein (1989), we distinguish between truly irrelevant information, that is, trivial information that is not diagnostic of any aspects of personality, and pseudo-relevant information, that is, information that is irrelevant to the judgment in question but that is diagnostic of other aspects of personality.

Phenomenon 10: Stereotypes affect impressions in the presence of truly irrelevant information. Truly irrelevant, trivial behaviors appear not to affect the impact of stereotypes on judgment. Thus, when targets of different gender, profession, or college major are known to have performed uninformative behaviors such as get a haircut, find 20 cents, arrive in class on time, or have a slice of pizza, the stereotypes associated with these targets do affect judgments about their stereotypic traits and trait-related behaviors (Hilton & Fein, 1989; Krueger & Rothbart, 1988; Locksley et al., 1980, Experiment 2; Macrae et al., 1992; Rasinski et al., 1985). Professional stereotypes also appear to affect impressions in the presence of minimal biographical information such as name, age, and address (Yzerbyt, Schardron, Leyens, & Rocher, 1994; significance not reported). Two out of the three of these studies that permit comparison of the effects of stereotypes in the presence of the truly irrelevant individuating information to their effects in its absence found that the effects of stereotypes were not diluted by the nondiagnostic individuating information (Hilton & Fein, 1989, Experiment 2; Locksley et al., 1980, Experiment 2). The third one did obtain

significant but relatively weak dilution (Hilton & Fein, 1989, Experiment 3).

Overall, then, it appears that the effects of stereotypes on impressions of targets are typically not mitigated by the presence of truly nondiagnostic, trivial individuating information about these targets. This highly intuitive finding requires little explanation: There is no reason to expect information to affect judgments to which it is completely irrelevant. In our IMP simulation, trivial information has no associations with any traits or trait-related behaviors and so has no impact on their activation levels.

Phenomenon 11: Nondiagnostic but pseudo-relevant information can eliminate or dilute the effects of stereotypes. It is also the case, however, that nondiagnostic information often eliminates or dilutes the effects of stereotypes on impressions. This pattern is typically found for pseudo-relevant information, that is, information that is nondiagnostic of the traits or behaviors in question but is diagnostic of other aspects of personality. Several studies have found that pseudo-relevant individuating information can completely eliminate the effects of stereotypes on judgments of the target. Thus a description of a student that included information about academic record, parents' occupations, and college friends, all of which were quite unremarkable for college students and unrelated to the dimension of self-control, sufficed to eliminate the effects of the stereotypes of night and day people on ratings of the student's self-control and other stereotypic traits (Locksley et al., 1982). Similarly, in the presence of comparable pseudo-relevant but nondiagnostic information, gender had no effect on ratings of targets on gender-related traits (Glick et al., 1988; Hilton & Fein, 1989, Experiment 1), and college majors had no effect on predictions of targets' shock tolerance and movie attendance (Nisbett et al., 1981, Experiment 3).

Several other studies have found that stereotypes do affect judgments about targets in the presence of such nondiagnostic but pseudo-relevant information, but typically these effects are diluted, that is, weaker than the effects of the stereotype in the absence of the individuating information. In the presence of nondiagnostic but pseudo-relevant information about targets, college majors affected predictions of targets' shock tolerance and movie attendance (Nisbett et al., 1981, Experiments 1 and 2) and ratings of target competitiveness (Hilton & Fein, 1989, Experiments 2 and 3), and gender affected ratings of the target's gender-related traits (Heilman, 1984) and suitability for gender-related jobs (Glick et al., 1988). Four out of the five of these studies that permit comparison of the effects of stereotypes in the presence of nondiagnostic but pseudo-relevant individuating information to their effects in its absence found that the effects of stereotypes were diluted by the nondiagnostic individuating information (Hilton & Fein, 1989, Experiments 2 and 3; Nisbett et al., 1981, Experiments 1 and 2).

Because there is little evidence bearing on the mechanisms underlying the dilution effect, our account of this effect is speculative. We propose that the diluting effect of pseudo-relevant information occurs because this information gives rise to the stereotype of *ordinary, average person*. Nondiagnostic information found to dilute or eliminate the effects of stereotypes typically portrays the individuals as bland, unremarkable, and aver-

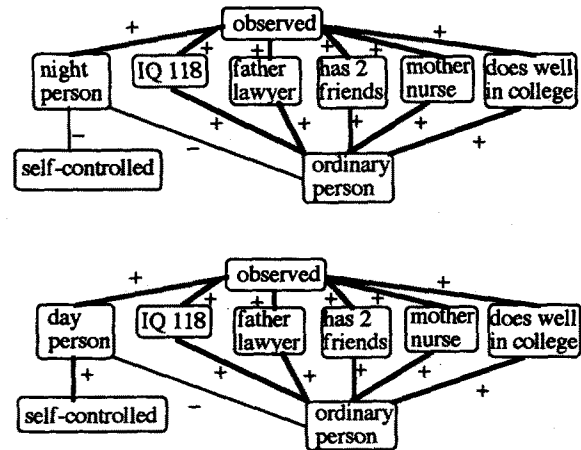


Figure 7. The effects of stereotypes are diluted in the presence of nondiagnostic individuating information (Phenomenon 11). The difference in activation levels of *self-controlled* under night person and day person decreases as the number of nondiagnostic pieces of information increases.

age. Consider, for example, the following information: "Gene W. has an IQ of 118 and does fairly well in his college classes. His mother works as a nurse and his father is a lawyer in the town where Gene grew up. Gene has a couple of friends in college, one of whom went to the same high school as he did." This information is nondiagnostic of self-control (Locksley et al., 1982). But it does bring to mind an average, typical college student. We assume that the stereotype of the typical, ordinary person inhibits all other stereotypes. This is because any other stereotype is noteworthy only inasmuch as it implies that members of the stereotyped group differ systematically from ordinary people.

As shown in Figure 7, the stereotype of *ordinary person* inhibits the observed stereotypes and so undermines their impact on judgment. In the absence of any individuating information, *self-controlled* has positive activation for a day person and negative activation for a night person, and the difference between these two activation levels (i.e., the effect of stereotypes) is substantial (.78). In the presence of the five pieces of individuating information depicted in Figure 7, this difference becomes much smaller (.18). Additional simulations indicated that, consistent with empirical findings (Nisbett et al., 1981), the effects of stereotypes decreased, that is, there was a greater dilution effect, as the number of items of nondiagnostic information increased. For example, the difference in the activation levels of *self-controlled* for day person and for night person was reduced to .42 in the presence of two items, and further reduced to .26 in the presence of four items. The successful simulation of these empirical findings suggests that our account is plausible. However, further empirical work is necessary to examine the veracity of our assumptions about the role of the stereotype of the ordinary person in producing the dilution effect.

Occasionally, stereotypes will affect impressions in the presence of pseudo-diagnostic information because the information becomes diagnostic when it characterizes members of stereo-

typed groups. Such stereotype-driven construal could change the meaning of the pseudo-diagnostic behaviors. For example, the behavior “recently bought the latest book of a best-selling author” appears nondiagnostic of aggression when described alone. But a construction worker said to engage in this behavior may be envisioned as buying a murder story, and a housewife as buying a romance. Thus construed, the behaviors appear differentially diagnostic of aggression. Therefore, the construction worker who bought a best-seller may be viewed as more aggressive than a housewife who bought a best-seller (Krueger & Rothbart, 1988) because the two are assumed to have bought different books (Kunda & Sherman-Williams, 1993). In a similar vein, Kunda and Sherman-Williams (1993) found that when a car salesman or a librarian were known to have engaged in a behavior that seems relatively nondiagnostic of extraversion, attending a party organized by their employer where they knew very few people, this information was construed quite differently as a function of the stereotypes, and the stereotype-driven construals were diagnostic of extraversion. The car salesman was assumed to have engaged in highly extraverted behaviors such as introducing himself to many people, whereas the librarian was assumed to have been relatively reserved. The effects of stereotypes on ratings of targets’ extraversion in the presence of the relatively nondiagnostic attendance at the party were mediated by the extraversion of subjects’ highly diagnostic construals of how the target behaved there.

To summarize, it appears that truly nondiagnostic, trivial behaviors have little impact on the effects of stereotypes on impressions of the persons they characterize. In contrast, pseudo-relevant information that is diagnostic of personality in general but not of the specific judgment subjects are asked to make tends to dilute or eliminate the effects of stereotypes. According to our model, stereotypes continue to affect judgment in the presence of truly nondiagnostic information because such information has no associations with the characteristics relevant to the judgment in question. And the effects of stereotypes are diluted or eliminated in the presence of pseudo-diagnostic information because this information gives rise to the stereotype of an ordinary person, which dampens the effects of stereotypes. On occasion, stereotypes may affect judgments of targets in the presence of either type of nondiagnostic information if different stereotypes cause this information to be construed as differentially diagnostic of these judgments.

Stereotypes and Other Stereotypes

Phenomenon 12: A stereotype can influence the activation of subtypes of another stereotype. In most research on the effects of stereotypes on impressions, participants are informed that a target individual belongs to a single stereotyped group—typically a particular race, gender, profession, or college major—and the effects of that group’s stereotype on impressions of the target are examined. Real people, however, belong to multiple social categories. Thus, we may encounter a White female doctor, a Black male political-science major, and so on. In some cases, perceivers may classify such individuals into preexisting subtypes. For example, Devine and Baker (1991) showed that people have subtypes for the combinations *Black businessman*

and *Black athlete*. And Brewer (1988) argued that people may have distinct subtypes for combinations of the most important social categories: gender, race, age, and profession. Thus Brewer assumes that people may have separate representations for young and old businessmen, for male and female doctors, and for many comparable combinations.

It is unlikely, however, that people have preexisting representations for all possible combinations of all stereotypes. When people encounter novel combinations of stereotypes, these may be integrated into coherent impressions of the target through the same processes used to integrate any other types of information. When an individual is known to belong to one stereotyped group, this knowledge may determine how one interprets information about other stereotyped groups to which that individual belongs. One stereotype may lead to the activation of particular subtypes of the other. Thus Kunda et al. (1990) reported that some subjects described a person known to be Harvard educated and a carpenter as an innovative master carpenter. And Hastie, Schroeder, and Weber (1990) reported that some subjects described a person known to be male and a nurse as a psychiatric nurse. According to our model, the relatively atypical subtype of carpenter, *master carpenter*, may have been activated because it had excitatory links with *Harvard educated* and its associates (e.g., *at top of profession*), whereas the more typical meanings of carpenter were inhibited by *Harvard educated*. This process could also account for the finding that the descriptions of people belonging to two stereotyped groups often contain emergent attributes, that is, attributes that are not used to describe members of either of the constituent stereotypes (Hastie et al., 1990; Kunda et al., 1990). This would occur if the atypical subtype activated is associated with attributes that are unique to it. For example, the Harvard-educated carpenter could be artistic, even though neither carpenters nor Harvard-educated people are viewed as artistic, if the activated subtype *master carpenter* is believed to be artistic.

Our model would also suggest that the many stereotypes applicable to an individual could jointly determine that individual’s probable traits and behaviors. When two stereotypes have different associations with a given trait, these associations may constrain each other. For example, a feminist is viewed as well educated, and a bank teller as not well educated. When a person is known to be a feminist and a bank teller, *well educated* will therefore be activated by *feminist* and inhibited by *bank teller*, and so its ultimate activation may be intermediate to those obtained when each of the stereotypes is presented alone. Kunda et al. (1990) report such averaging for this particular example as well as for a substantial proportion of the combinations they examined. IMP easily simulates such results.

Phenomenon 13: Combinations of stereotypes and traits can be creative. We described how some of the findings on how people combine social stereotypes may result from automatic processes. There is reason to believe, however, that people often turn to more elaborate causal reasoning when attempting to reconcile how a person could belong to two different groups whose stereotypes have conflicting implications. Kunda et al. (1990) found that many participants used causal reasoning in their open-ended descriptions of persons belonging to conflicting stereotypes. For example, a person who was gay and a con-

struction worker was described as "compensating for his 'unusual' tendencies by taking a job that is considered 'manly.'" Similar examples were reported by Hastie et al. (1990). Kunda et al. (1990) also found that subjects were more likely to generate such causal accounts the more surprising they found the combination, suggesting that causal reasoning was triggered by surprise.

Research on how people combine social stereotypes has also found many examples of emergent attributes, that is, attributes that are used to characterize a person belonging to multiple stereotypes even though they are not used to describe persons belonging to any one of the constituent stereotypes (Hastie et al., 1990; Kunda et al., 1990). Emergent attributes may result from automatic processes, as described earlier; but they may also result from causal reasoning. For example, a person described as Harvard educated and a carpenter was viewed as more likely to be nonmaterialistic than was a person belonging to each of those categories alone. This attribute seems likely to be the product of a causal analysis—perceivers assume that since a Harvard-educated person can choose more profitable careers, this person must be nonmaterialistic to have chosen to become a carpenter. Participants reported many examples of such explicit causal analyses (Kunda et al., 1990). It is often difficult to tell whether a particular emergent attribute resulted from this kind of causal reasoning or from more automatic spreading activation in which an unusual subtype of one stereotype is boosted by the other stereotype (in this example, one may evoke the subtype of the Harvard-educated hippie, who is believed to be nonmaterialistic). But it seems likely that both types of processes contribute to the combination of stereotypes.

Controlled, inferential processes are similarly implicated in research on how people combine different personality traits to form a coherent impression. Subjects asked to describe a person characterized by conflicting traits often resolve the conflicts through causal reasoning (Asch & Zukier, 1984). Asch and Zukier outlined several modes used by subjects to resolve conflict among traits. Although some of these modes may imply automatic processes, several common modes of resolution clearly involve causal analysis. These include cause-effect reasoning, in which one trait is viewed as the reason for the other (e.g., "dependence often breeds hostility") and means-end analysis, in which one trait is viewed as a means for accomplishing the other (e.g., "to be kind it is sometimes necessary to be strict"). Comparable causal reasoning has also been used by subjects describing persons characterized by congruent traits (Casselden & Hampson, 1990). In these studies, the causal reasoning may have been triggered not only by surprise but also by explicit instructions to describe "how these qualities may be related" (Asch & Zukier, 1984) or how they "go together in a person" (Casselden & Hampson, 1990), which appear to encourage causal reasoning.

IMP does not model this kind of causal reasoning. However, we believe that the products of causal reasoning may be integrated with and constrained by other information known to characterize the individual, through the processes modeled by IMP.

Order Effects

All of the phenomena discussed so far focus on how perceivers who have no prior expectations about the target persons form

impressions of these persons after observing several of their characteristics simultaneously. Many everyday encounters with strangers have this format. For example, on meeting a new person, one may notice simultaneously that she is White, female, attractive, well dressed, and nervous. We have concentrated on the impressions formed in such encounters with strangers because they have been the focus of most of the research on how stereotypes and individuating information affect impressions.

But in everyday life people do not interact only with strangers and do not always obtain all their information about a given person at the same time. One often has prior expectations about others. One may observe a person doing one thing today and something else tomorrow. One may notice or be told about a person's gender or race before one has had an opportunity to observe this person's behavior. Or one may observe a person's behavior before finding out his or her profession. There is reason to believe that the order in which information about a person is observed will affect one's impression of this person. In a classic experiment, Asch (1946) showed that a person will be viewed as more likable when described as "intelligent-industrious-impulsive-critical-stubborn-envious" than when described by the very same list of traits presented in the opposite order (see also Jones & Goethals, 1972). In Asch's view, this primacy effect occurred because some of the same traits gain a positive meaning when preceded by a positive trait such as *intelligent* but gain a negative meaning when preceded by a negative trait such as *envious*. For similar reasons, stereotypes may have greater impact on impressions when observed before rather than after individuating information has been observed (e.g., Bodenhausen, 1988).

IMP can readily model such primacy effects. To simulate order effects, we first run IMP with only the first piece of observed information and its associates. For example, to simulate Asch's (1946) findings, the network is initially given only the information that the person is intelligent. *Intelligent* is strongly associated with traits such as *sensible* that make a person likable, and weakly associated with traits such as *arrogant* that make a person dislikable. After the network has settled, we run it again, adding another piece of observed information, *stubborn*. This second run differs importantly from all the simulations we have presented so far. In our previous runs, all units start off with the same activation of close to zero because the perceiver is assumed to know nothing about the target to begin with. In contrast, in this second run, the units start off with the activation values on which the network has settled after the first run, that is, after observing the first piece of information about the target. Therefore, because *likable* has become strongly activated following the first run, it can now color the interpretation of *stubborn*. *Stubborn* is strongly associated with dislikable traits such as *bullheaded* and weakly associated with more likable traits such as *persevering*. Because the second run began with *likable* strongly activated, the more likable aspects of *stubborn* will be more strongly activated than they would have been if *stubborn* were presented alone. Thus we find that if IMP is given first the information that the person is intelligent and then the information that the person is stubborn, *likable* ends up with stronger activation than it does when IMP is given the same observed traits in the opposite order. For the sake of simplicity, our exam-

ple focused on only two characteristics that are observed in different orders. The same procedure of sequential runs, each beginning with the values obtained from the previous one, can be used to model successive observations of any number of characteristics.

Summary

We have shown that our parallel-constraint-satisfaction theory of impression formation can account for most findings reported in the stereotype literature. Our review of this literature led to the following conclusions: Stereotypes can affect the meaning of behavior and traits, and these can affect the meaning of stereotypes. Stereotypes affect ratings of individuals in the absence of individuating information. Their impact in the presence of individuating information depends on the ambiguity of this information and on the judgment task. Stereotypes typically do not affect trait ratings in the presence of unambiguous diagnostic information but do affect trait ratings in the presence of ambiguous diagnostic information by affecting the construal of this information. Stereotypes do affect trait-related behavioral predictions even when their impact on the traits has been undermined by individuating information because of the associations of the predicted behaviors with other stereotypic traits. Stereotypes do affect impressions in the presence of trivial non-diagnostic information, but their impact is diluted or eliminated in the presence of pseudo-diagnostic individuating information that has implications for other aspects of the target's personality. And multiple stereotypes can jointly affect impressions of a person, by affecting each other's meaning. All of these findings can be modeled by IMP's constraint networks, which suggests that these phenomena may be due to relatively automatic processes that require little conscious inference.

We also pointed to a small number of phenomena that do seem to imply more controlled processes involving conscious inference and elaborate causal reasoning. These include creative combinations of surprising pairs of stereotypes or traits, and the creation of novel subtypes to account for individuals whose attributes strongly violate their group's stereotype. The phenomena that require causal reasoning appear to occur when perceivers are confronted with particularly surprising combinations of information (Kunda et al., 1990). We also suggested that certain tasks require causal reasoning. These include assessment of the causes underlying behavior and assessment of guilt in the presence of uncertain evidence. Controlled inference may also be triggered by instructions that emphasize reasoning and by motivation.

Comparison With Serial Theories

Our parallel processing model of impression formation assumes that stereotypes and individuating information constrain each other's meaning and jointly influence impressions. These ideas are not entirely incompatible with the serial models developed by Brewer (1988) and by Fiske and Neberg (1990). Both serial models posit that, under some circumstances, stereotypes and individuating information will jointly contribute to impressions. In Brewer's model, such joint influence occurs when in-

formation about the target is personalized and perceivers engage in bottom-up processes. When such personalization occurs, the stereotype is represented as but one feature of the target, to be integrated with other target information. In Fiske and Neberg's model, stereotypes and individuating information will both influence impressions, albeit to different degrees, when perceivers engage in processes that fall in the midrange of their hypothesized continuum. The notion that the integration of stereotypes with individuating information occurs through parallel processes is compatible with both Brewer's and Fiske and Neberg's models. Indeed, Fiske and Neberg (1990) raise this possibility.

Although our model and the serial ones may agree on how stereotypes are integrated with individuating information, they clearly differ in their assumptions about when such integration will take place. We assume that whenever stereotypes and individuating information are both observed, they will jointly influence impressions. In contrast, both Brewer's (1988) and Fiske and Neberg's (1990) models assume that processes in which stereotypes are integrated with individuating information occur, if at all, only after stereotype-dominated processes have failed to account for the information at hand or failed to satisfy one's goals. Thus, in our model, stereotypes are given no special status and are treated no differently from any other type of information. But in the serial models stereotypes play a dominant role in impression formation, in that stereotype-driven processes take priority over more personalized, individuated ones. Our model does allow stereotypes to dominate impressions when they are observed before individuating information, through primacy effects. But it also allows individuating information to dominate impressions in the same way when it is observed first, and assumes that both kinds of information have equal status when they are observed simultaneously. We also suggest that stereotypes will appear to dominate impressions when they have stronger or more numerous associations with the dependent measures than does the individuating information and when they are less open to multiple construal than the individuating information. However, we maintain that individuating information will also appear to dominate impressions when the reverse conditions hold true.

Our parallel-constraint-satisfaction theory and computational model have several advantages over the serial models proposed by Brewer (1988) and by Fiske and Neberg (1990). Our theory has greater explanatory breadth in that it can account for the full range of findings concerning how different sources of information about a person are influenced by each other and are integrated to form a coherent impression. Our account also possesses greater precision and simplicity and avoids the need for the potentially problematic distinction between stereotypes and individuating information. We now elaborate on each of these advantages.

Explanatory Breadth

There are several phenomena that are more readily explained by our theory than by the alternative, serial ones. The most important of these is the finding that the effects of stereotypes depend on the nature of the judgment task: Clear diagnostic indi-

viduating information undermines the effects of stereotypes on ratings of the target's traits but does not undermine the effects of stereotypes on predictions about the same target's future trait-related behavior (Kunda, Sinclair, & Griffin, 1995). This finding is problematic for the serial models, which assume that the key to understanding the effects of stereotypes and individuating information on impressions lies in identifying the appropriate kind of processing that the perceiver engages in. In those models, stereotypes, individuating information, and goals determine whether the target will be processed through category-based, attribute-based, or intermediate processes. The models imply that the kind of process engaged in should determine the effects of the different kinds of information on all dependent measures. When category-based processes prevail, the stereotype should affect both trait ratings and behavioral predictions. And when individuating processes prevail, the stereotype should affect neither; there is no reason to distinguish among measures. In contrast, the differential effects of stereotypes on traits and on behavioral predictions fall naturally out of our model as a function of the patterns of associations among traits, behaviors, and the different aspects of stereotypes (see Figure 6).

The serial models, which assume that one stereotype prevails at a time, also have trouble accounting for findings showing that multiple stereotypes can jointly affect impressions. This is a serious problem, because most individuals belong to multiple categories, and perceivers base their impressions on more than one category (Hastie et al., 1990; Kunda et al., 1990). Brewer's (1988) model partially addresses this issue by assuming preexisting representations of combinations of the most important stereotypes. And indeed, there is reason to believe that perceivers often do categorize individuals at the level of preexisting subtypes, such as Black woman, that combine race and gender (Stangor, Lynch, Duan, & Glass, 1992). Fiske and Neuberg (1990) assume that one of the many social categories that a person belongs to will be selected to function as the dominant stereotype, and describe the factors affecting this selection. But both Brewer's and Fiske and Neuberg's models have trouble accounting for how novel combinations of stereotypes can jointly influence impressions. In contrast, our model can integrate multiple stereotypes much as it integrates any other sources of information.

Simplicity

The serial models assume that the diverse phenomena that we have discussed each arise from a different kind of process occurring at different stages of impression formation. For example, Fiske and Neuberg (1990) postulated that stereotypes affect the meaning of attributes through confirmatory categorization of these attributes; behavior-driven recruitment of atypical subtypes of a stereotype results from recategorization, after a failed initial categorization; and integration of attributes occurs through piecemeal processes, after both initial categorization and recategorization have failed. Brewer (1988) also assumed that these phenomena result from different kinds of processes. In contrast, we parsimoniously account for all of these phenomena through the same basic process of parallel constraint satisfaction.

Precision

Earlier models remained vague about the structure and processes affecting impression formation. In contrast, we have spelled out the mechanisms that we propose in sufficient detail that they can be simulated by a running computer program. The program IMP provides not only qualitative simulations of the major phenomena concerning impression formation but also quantitative output that can be compared with the results of psychological experiments (see, for example, Tables 1-3).

Avoiding the Distinction Between Stereotypes and Individuating Information

The final advantage we claim for our theory of impression formation is that it does not require the problematic distinction between stereotypes and individuating information that plays a central role in Brewer's (1988) and in Fiske and Neuberg's (1990) models. In both models stereotypes dominate impressions, and individuating information is used only to help categorize or recategorize the target unless there is strong pressure to form an individuated impression, in which case the stereotypes play a subsidiary role. But it is not always clear which kinds of information constitute a stereotype and which constitute individuating information. For example, would one want to consider terms such as sociable, brainy, or athletic to be individuating information or stereotypes? All of these terms are traits, which are typically considered to be individuating information, but in what way do they differ from stereotypes? Structurally, these and many other trait terms seem comparable to stereotypes. Both involve rich networks of associations including other traits as well as behavioral and physical expectations. One may imagine a sociable or a brainy type of person much as one imagines a lawyer or a feminist. As with many nonsocial concepts, both traits and stereotypes can be organized hierarchically (Cantor & Mischel, 1979; Hampson, John, & Goldberg, 1986). Some broad and abstract superordinate traits such as extraverted or introverted may be less rich and differentiated than some narrow stereotypes such as comedian or politician (Andersen & Klatzky, 1987). But traits, like stereotypes, vary in breadth (Hampson et al., 1986), and some less abstract subordinate traits such as sociable or brainy may be more rich and differentiated than some broad superordinate stereotypes like woman or American.

It is also difficult to distinguish among traits and stereotypes on the basis of reference class. Stereotypes, it may be argued, refer to knowledge structures about human social groups (Hamilton & Sherman, 1994). But what constitutes a social group? Demographics, professions, and roles are typically used as the basis for grouping. But it may be just as meaningful to group people according to their intelligence or sociability, attributes that are typically viewed as individuating information. Thus many attributes that are typically viewed as individuating information appear indistinguishable from stereotypes both structurally and in terms of reference class.

The distinction becomes even more problematic when one considers that in real life, unlike in experiments, perceivers often do not observe social categories directly. Rather, these need

to be inferred from the target's attributes. Thus, one does not observe a wealthy, educated, Black comedian. Instead, one observes a person with a certain skin color, clothes, accent, vocabulary, behavior, and many other attributes from which the person's membership in various social categories may be inferred. If one assumes that such complex arrays of features will give rise to only one category that will then dominate impressions, one must also assume that the same features can count sometimes as stereotypes and sometimes as individuating information. Thus the attribute *well dressed* may trigger social class on some occasions, and so function as a stereotype, but may function as individuating information on other occasions. In a similar way, the behavior *tells jokes* may either trigger the stereotype of comedian or function as individuating information. The distinction between stereotypes and individuating information thus becomes based on the role that the information plays in a given impression formation task rather than on any inherent structural differences between the two classes of information. The serial models proposed by Brewer (1988) and Fiske and Neberg (1990) assume that although some features, such as gender, race, or social roles, are more likely than others to serve as the dominant stereotype, the selection of the dominant stereotype also depends on transient factors such as the feature's accessibility and contextual salience and the perceiver's mood and goals. Fiske and Neberg (1990) ultimately resolve the problem of definition in a somewhat circular manner, by concluding that a feature constitutes a category if it functions as a category, that is, dominates impressions. This solution, however, robs the distinction between stereotypes and individuating information of much of the meaning and connotations that it typically carries.

The distinction between stereotypes and individuating information is not an issue from our perspective, because we give the two classes of information equal status. The effects of information on impressions do not depend on whether this information constitutes a stereotype or individuating information but on its unique pattern of associations with other information.

Do Stereotypes Dominate Impressions?

Unlike the serial models, our account does not assume that stereotypes dominate impressions, or that they are used earlier and more automatically than are other types of information. Because the view that stereotypes dominate impressions has gained considerable acceptance (Hamilton & Sherman, 1994), we argue next that the evidence that stereotypes dominate impressions is inconclusive at best, and that there are reasons to believe otherwise. We begin by discussing the three major arguments for the dominance of stereotypes and then present evidence that stereotypes do not dominate impressions in the presence of individuating information.

Evaluating the Arguments for the Dominance of Stereotypes

1. *Stereotypes are pervasive.* Much of the argument for the dominance of stereotypes rests on their pervasive impact on impressions and on the interpretation of individuating informa-

tion (Fiske & Neberg, 1990). We have shown, however, that we can account for the full range of stereotypic effects without assuming that stereotypes dominate impressions.

2. *Stereotypes serve as heuristics.* It has been suggested that stereotypes function as heuristics in that they are applied automatically and require few cognitive resources. The use of individuating information, in contrast, is more effortful, and so this information is used only when perceivers have sufficient resources and are motivated to form accurate impressions. This argument rests on the claim that stereotypes are more likely to dominate impressions in the presence of individuating information when perceivers have limited cognitive resources. Evidence marshaled to support this view consists of the findings that stereotypes have greater impact on impressions when resources are limited because perceivers are not at their optimal time of day (Bodenhausen, 1990), because they are cognitively busy (Gilbert & Hixon, 1991; Pendry & Macrae, 1994), or because they are happy (Bodenhausen, Kramer et al., 1994) or angry (Bodenhausen, Sheppard et al., 1994), or are under time pressure (Kruglanski & Freund, 1983). And the impact of stereotypes decreases when perceivers are required to expend more effort on the judgment, because it is complex (Bodenhausen & Lichtenstein, 1987) or because they expect to be accountable for their judgments (Bodenhausen, Kramer et al., 1994).

If the effects of individuating information are indeed more likely than the effects of stereotypes to be disrupted when cognitive resources are depleted, this would imply that the use of stereotypes is more automatic and less effortful than the use of individuating information. However, to support this conclusion it is necessary to show not only that reduction in cognitive resources increases reliance on stereotypes, as the studies cited above have done, but also that it decreases reliance on individuating information. For example, it would be necessary to vary targets' race and ability orthogonally and show that conditions that increase the impact of the target's race on impressions also decrease the impact of the target's ability on impressions. The demonstration that cognitive load causes such disruption of the impact of individuating information on impressions is crucial to the argument that use of individuating information requires more controlled processes than does use of stereotypes. Unfortunately, none of these studies varied individuating information. Instead, they all relied on a single description of a target, which was presented as associated with different stereotypes. It therefore remains quite possible that in these studies cognitive load, which was shown to increase reliance on stereotypes, may have also, at the same time, increased reliance on individuating information. One study that varied both individuating information and stereotypes did not support the view that the processing of individuating information requires more resources than the processing of stereotypes (Pratto & Bargh, 1991). This study, which created cognitive load by speeding up the presentation of information, found that load eliminated the impact of stereotypes on global trait ratings but did not affect the impact of individuating information on these ratings. This study suggests, therefore, that, if anything, stereotypic processing requires more resources than the processing of individuating information. It does not make this point conclusively, however, because other, more specific trait measures showed that load

reduced the impact of both stereotypes and individuating information.

It has also been shown that cognitive load may disrupt the activation of stereotypes, which suggests that stereotypes are not always activated automatically and effortlessly (Gilbert & Hixon, 1991). This study could not assess, however, whether cognitive load also undermined the effects of stereotypes on impressions of targets in the presence of individuating information. This is because, as is typically the case, stereotypes had no such effect on impressions in the first place, in the absence of cognitive load.

In sum, there is no conclusive support for the notion that stereotypes are used more automatically than individuating information. Cognitive load increases the impact of stereotypes on impressions, but there is no evidence that it reduces the impact of individuating information. And cognitive load has also been shown to decrease the activation of stereotypes.

3. *Stereotypes organize knowledge under high information load.* Several studies suggest that when perceivers are exposed simultaneously to multiple unfamiliar target persons who are each characterized by multiple attributes, they tend to organize this information categorically rather than on a person-by-person basis (Brewer, 1988). Thus, unfamiliar names do not serve to organize unconnected pieces of information associated with them (although familiar names do), implying lack of person-based organization (Pryor & Ostrom, 1981). And comments made by members of a given social category are more likely to be mistakenly attributed to another member of that category than to a member of a different category, suggesting category-based organization (S. E. Taylor, Fiske, Etcoff, & Ruderman, 1978). Such findings show that when perceivers are confronted with multiple novel target persons each of whom is characterized by several disjoint and not very memorable pieces of information, they find it difficult to recall precisely which piece of information was associated with which person. Under such circumstances, existing knowledge structures such as social categories may better serve to organize the mass of incoming information, if only because they serve to reduce the number of organizational units from a relatively large number of persons to a more manageable number of categories that can be handled in short-term memory (Miller, 1956). But such findings do not imply that social categories will dominate impressions of individuals who are encountered one at a time. Indeed, perceivers appear to engage in different processes when thinking about persons encountered individually versus in groups (Brewer, 1988).

Thus it appears that there is little conclusive support for the notion that stereotypes dominate impressions.

Comparing the Impact of Stereotypes and Individuating Information on Impressions

One way of determining the relative impact of stereotypes and individuating information on impressions is by examining studies in which both types of information were varied orthogonally. If stereotypes dominate impressions they should exert greater impact on impressions than does individuating information. But this is clearly not the case. Rather, when both types

of information were varied in the same experiment, the effects of stereotypes were typically dwarfed in comparison to the effects of individuating information. We conducted a meta-analysis to determine the magnitudes of the effects of stereotypes and individuating information on impressions. We have identified 28 articles reporting 40 studies in which individuating information and stereotypes were varied orthogonally in the same study. In all but one of these studies, individuating information had larger effects on impressions than did stereotypes. The average effect size for individuating information in these studies ($r = .69$) was substantially larger than the average effect size for stereotypes ($r = .19$). Put differently, on average, individuating information accounts for 48% of the variance in target evaluations in these studies, whereas stereotypes account for only 4% of that variance.⁴ It is difficult to see how one can maintain the view that stereotypes dominate impressions in the face of such findings.

Future Directions

We have modeled situations in which different kinds of information about a person—group membership, behavior, traits, and so on—have been observed, and we showed how these pieces of information can be shaped by each other and integrated to form an impression. There are several important ways in which this model can be expanded in the future.

Integrating Automatic and Controlled Processes

We view impression formation as resulting from a complex interplay between controlled and automatic processes. We have focused here on the more automatic aspects of impression formation. Other models have focused on the more controlled ones such as the generation and evaluation of hypotheses (Read & Miller, 1993; Thagard, 1992). The literature on text comprehension, which influenced our thinking, is similarly bifurcated, with some models focusing predominantly on the relatively automatic processes of word disambiguation (Kintsch, 1988), and others on the more controlled processes of inference construction (Graesser, Singer, & Trabasso, 1994). Ultimately, a comprehensive model, encompassing both types of processes and spelling out their interactions, should be developed.

Stereotype Activation

Our accounts of the empirical phenomena, like the research on which they were based, assume that a stereotype is commu-

⁴ Articles used in this analysis are starred in the reference list. We followed the procedures described in Footnote 3 for locating studies and calculating results. We included all studies that varied both individuating information and stereotypes in the same study and that reported dependent measures assessing impressions of a specific target person (be they traits, behavioral predictions, or attributions). When a study reported more than one measure of impressions of target, these measures were averaged so that each study yielded a single effect size for stereotypes and a single effect size for individuating information. Results reported separately for separate groups of subjects were treated as independent studies.

nicated explicitly to the perceiver. Some stereotyped categories such as gender and race are observed practically without inference, but others require greater inference and may be held more tentatively. Such inferred stereotypes may include profession, political and religious affiliation, and social types such as class clown, fraternity member, or mother. There is very little research examining how such stereotypes are inferred. Brewer (1988) reported some preliminary work exploring which stereotypes are most likely to be triggered automatically, but we are aware of no research examining which stereotypes will be inferred from complex arrays of information comparable to those encountered by naive observers in their daily lives. As well, it is unlikely that all of the possible stereotypes that can characterize a given person will be activated simultaneously. A given person may be a woman, a Jew, an American, middle-aged, a professor, a mother, a wife, a baseball fan, and the list could go on. Typically, only a subset of these pertinent stereotypes will be activated. There is evidence that contextual salience can determine the activation of racial and gender stereotypes (S. E. Taylor, 1981), but there is no research on how one chooses among stereotypes that are less visually obvious. From the perspective of our parallel-constraint-satisfaction theory of impression formation, inference to stereotypes is cognitively identical to the kind of inference from stereotypes that we have so far modeled. Our program IMP is capable without modification of simulating how observations of traits and behaviors can lead to the activation of stereotypes. Whether this is a psychologically accurate account of inference to stereotypes can only be determined by future psychological experiments.

Motivation

Motivation can play an important role in determining the extent as well as the direction of processing. The motivation to be accurate may lead to deeper processing and to greater reliance on controlled processes (Fiske & Neuberg, 1990), and the motivation to arrive at particular conclusions may affect impressions of targets as well as one's stereotypes (Klein & Kunda, 1992; Kunda, 1990). As the relation between automatic and controlled processes is investigated, the effects of accuracy goals on the two kinds of processes and their interaction should be explored. It may be possible to incorporate directional goals into the model by assuming that they increase the activation of desired constructs and thereby constrain impressions.

Affect

Emotions and affect may also play an important role in the representation of stereotypes (Esses, Haddock, & Zanna, 1993) and influence their use in impression formation (Bodenhausen, Kramer, & Susser, 1994; Bodenhausen, Sheppard, & Kramer, 1994; Jussim et al., 1995). This, too, should be incorporated into future models. One way of incorporating affect could be by adding affective units to the knowledge representation and treating them as comparable in status to cognitive units. For example, stereotypes and behaviors may activate or deactivate affective units such as *like*, *scary*, or *exciting* much like they activate and deactivate associated cognitive units such as traits

or behaviors. Mischel and Shoda's (1995) model of personality incorporates affective units in this manner. Our model could include such affective units without modification. By conceptualizing affect in this manner we could account for the role that affect plays in guiding impressions of stereotyped people (Jussim et al., 1995) and attitudes towards them (Esses et al., 1993). However, future models may need to be modified if affect is shown to influence not only impressions of targets but also the nature of the process through which these impressions are formed.

Stereotype Change

Our simulations have taken the knowledge base, including units and the weights between them, as given. They do not address the question of how incoming information may alter one's knowledge about stereotypes, behaviors, and their associations. Our IMP program uses a "local" connectionist representation in which each unit has an identifiable interpretation in terms of a particular characteristic. In contrast, many connectionist models acquire distributed representations that contain noninterpreted hidden units whose weights to input and output units are learned by training on numerous examples (Rumelhart & McClelland, 1986). When stereotypes do change, this may occur in such an incremental fashion. Stereotypes may also change through more elaborate causal reasoning, much like scientific theories change (Thagard, 1992; Weber & Crocker, 1983).

Other Domains of Application

In this article we focused on how our parallel-constraint-satisfaction model can explain how various types of knowledge about a person are integrated to form an impression of that person. But our model is general enough to be applicable to any judgment task that requires the integration of many sources of information. Peoples' self-concepts, their attitudes toward objects, people, and issues, their understanding of social situations, their intentions to engage in various behaviors, their perceptions of justice and fairness, and many other aspects of personality and social intelligence could be viewed as resulting from the integration of numerous pieces of information (cf. Mischel & Shoda, 1995). Parallel-constraint-satisfaction models provide new analytic tools for theorizing about these issues and can shed new light on many classic areas of social-psychological investigation.

References

- References marked with an asterisk indicate studies included in the meta-analysis.
- Andersen, S. A., & Klatzky, R. L. (1987). Traits and social stereotypes: Levels of categorization in person perception. *Journal of Personality and Social Psychology*, 53, 235-246.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, N. H. (1968). A simple model for information integration. In R. P. Abelson, E. Aronson, W. J. McGuire, T. M. Newcomb, M. J.

- Rosenberg, & P. H. Tannenbaum (Eds.), *Theories of cognitive consistency: A sourcebook* (pp. 731-743). Chicago: Rand McNally.
- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, 41, 303-314.
- Asch, S. (1987). *Social psychology*. Oxford: Oxford University Press. (Original work published 1952).
- Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology*, 46, 1230-1240.
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. A. Bargh & J. S. Uleman (Eds.), *Unintended thought* (pp. 3-51). New York: Guilford Press.
- Beckett, N. E., & Park, B. (1995). Use of category versus individuating information: Making base rates salient. *Personality and Social Psychology Bulletin*, 21, 21-31.
- Berndt, T. J., & Heller, K. A. (1986). Gender stereotypes and social inferences: A developmental study. *Journal of Personality and Social Psychology*, 50, 889-898.
- *Biernat, M., & Manis, M. (1994). Shifting standards and stereotype-based judgments. *Journal of Personality and Social Psychology*, 66, 5-20.
- *Bigoness, W. J. (1976). Effect of applicant's sex, race, and performance on employers' performance ratings: Some additional findings. *Journal of Applied Psychology*, 61, 80-84.
- *Bodenhausen, G. V. (1988). Stereotypic biases in social decision making and memory: Testing process models of stereotypic 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., Kramer, G. P., & Susser, K. (1994). Happiness and stereotypic thinking in social judgment. *Journal of Personality and Social Psychology*, 66, 621-632.
- *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., Sheppard, L. A., & Kramer, G. P. (1994). Negative affect and social judgment: The differential impact of anger and sadness. *European Journal of Social Psychology*, 24, 45-62.
- Bodenhausen, G. V., & Wyer, R. S. (1985). Effects of stereotypes on decision making and information processing strategies. *Journal of Personality and Social Psychology*, 48, 267-282.
- Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition* (pp. 1-36). Hillsdale, NJ: Erlbaum.
- Brewer, M. B., Dull, V., & Lui, L. (1981). Perceptions of the elderly: Stereotypes as prototypes. *Journal of Personality and Social Psychology*, 41, 656-670.
- Cantor, N., & Mischel, W. (1979). Prototypes in person perception. In L. Berkowitz (Eds.), *Advances in experimental social psychology* (pp. 3-53). New York: Academic Press.
- Casselden, A. P., & Hampson, S. E. (1990). Forming impressions from incongruent traits. *Journal of Personality and Social Psychology*, 59, 353-362.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82, 407-428.
- *Costrich, N., Feinstein, J., Kidder, L., Marecek, J., & Pascale, L. (1975). When stereotypes hurt: Three studies of penalties for sex-role reversals. *Journal of Experimental Social Psychology*, 11, 520-530.
- 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., & Emswiller, T. (1974). Explanations of successful performance on sex-linked tasks: What is skill for the male is luck for the female. *Journal of Personality and Social Psychology*, 29, 80-85.
- *Deaux, K., & Lewis, L. L. (1984). Structure of gender stereotypes: Interrelationships among components and gender label. *Journal of Personality and Social Psychology*, 46, 991-1004.
- *Deaux, K., & Taynor, J. (1973). Evaluation of male and female ability: Bias works two ways. *Psychological Reports*, 32, 261-262.
- *Del Boca, F. K., & Ashmore, R. D. (1980). Sex stereotypes and implicit personality theory: II. A trait-inference approach to the assessment of sex stereotypes. *Sex Roles*, 6, 519-535.
- 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.
- *Dipboye, R. L., Fromkin, H. L., & Wiback, K. (1975). Relative importance of applicant sex, attractiveness, and scholastic standing in the evaluation of job applicant résumés. *Journal of Applied Psychology*, 60, 39-43.
- *Dipboye, R. L., & Wiley, J. W. (1977). Reactions of college recruiters to interviewee sex and self-presentation style. *Journal of Vocational Behavior*, 10, 1-12.
- 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., 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: Academic Press.
- Feather, N. T., & Simon, J. G. (1975). Reactions to male and female success and failure in sex-linked occupations: Impressions of personality, causal attributions, and perceived likelihood of different consequences. *Journal of Personality and Social Psychology*, 31, 20-31.
- *Feldman-Summers, S., & Kiesler, S. B. (1974). Those who are number two try harder: The effect of sex on attributions of causality. *Journal of Personality and Social Psychology*, 30, 846-855.
- Fiske, S. T. (1990). Compare and contrast: Brewer's dual process model and Fiske et al.'s continuum model. In T. K. Srull & R. S. Wyer (Eds.), *Advances in social cognition*. Hillsdale, NJ: Erlbaum.
- 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. Zanna (Ed.), *Advances in experimental social psychology* (pp. 1-74). San Diego, CA: 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., & Von Hendy, H. M. (1992). Personality feedback and situational norms can control stereotyping processes. *Journal of Personality and Social Psychology*, 62, 577-596.
- Galambos, J. A., Abelson, R. P., & Black, J. B. (1986). *Knowledge structures*. Hillsdale, NJ: Erlbaum.
- Gilbert, D. T. (1989). Thinking lightly about others: Automatic components of the social inference process. In J. A. Bargh & J. S. Uleman (Eds.), *Unintended thought* (pp. 189-211). New York: Guilford Press.
- 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, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive busy-

- ness: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology*, 54, 733–740.
- *Glick, P., Zion, C., & Nelson, C. (1988). What mediates sex discrimination in hiring decisions? *Journal of Personality and Social Psychology*, 55, 178–186.
- Graesser, A. C., Singer, M., & Trabasso, T. (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101, 371–395.
- *Grant, P. R., & Holmes, J. G. (1981). The integration of implicit personality theory schemas and stereotype images. *Social Psychology Quarterly*, 44, 107–115.
- Griffin, D. W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. In M. Zanna (Ed.), *Advances in experimental social psychology* (pp. 319–359). San Diego, CA: Academic Press.
- Hamilton, D. L., & Sherman, J. W. (1994). Stereotypes. In J. R. S. Wyer & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 2, pp. 1–68). Hillsdale, NJ: Erlbaum.
- Hamilton, D. L., & Zanna, M. P. (1974). Context effects in impression formation: Changes in connotative meaning. *Journal of Personality and Social Psychology*, 29, 649–654.
- Hampson, S. E., John, O. P., & Goldberg, L. R. (1986). Category breadth and hierarchical structure in personality: Studies of asymmetries in judgments of trait implication. *Journal of Personality and Social Psychology*, 51, 37–54.
- *Hartman, S., Griffith, R. W., Crino, M. D., & Harris, O. J. (1991). Gender-based influences: The promotion recommendation. *Sex Roles*, 25, 285–300.
- Hastie, R., Schroeder, C., & Weber, R. (1990). Creating complex social conjunction categories from simple categories. *Bulletin of the Psychonomic Society*, 28, 242–247.
- Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley.
- Heilman, M. E. (1984). Information as a deterrent against sex discrimination: The effects of applicant sex and information type on preliminary employment decisions. *Organization Behavior and Human Performance*, 33, 174–186.
- *Heilman, M. E., & Stopeck, M. H. (1985). Attractiveness and corporate success: Different causal attributions for males and females. *Journal of Applied Psychology*, 70, 379–388.
- *Heneman, H. G. (1977). Impact of test information and applicant sex on applicant evaluation in a selection simulation. *Journal of Applied Psychology*, 62, 524–526.
- Hewstone, M. (1994). Revision and change of stereotypic beliefs: In search of the elusive subtyping model. *European Review of Social Psychology*, 5, 69–108.
- Hewstone, M., & Ward, C. (1985). Ethnocentrism and causal attribution in Southeast Asia. *Journal of Personality and Social Psychology*, 48, 614–623.
- 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. 69–121). Hillsdale, NJ: Erlbaum.
- Hilton, J. L., & Fein, S. (1989). The role of typical diagnosticity in stereotype-based judgments. *Journal of Personality and Social Psychology*, 57, 201–211.
- Holyoak, K. J., & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology*, 44, 265–315.
- Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, 13, 295–355.
- Holyoak, K. J., & Thagard, P. (1995). *Mental leaps: Analogy in creative thought*. Cambridge, MA: MIT Press/Bradford Books.
- *Jackson, L. A., Sullivan, L. A., & Hodge, C. N. (1993). Stereotype effects of attributions, predictions, and evaluations: No two social judgments are quite alike. *Journal of Personality and Social Psychology*, 65, 69–84.
- Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, 88, 67–85.
- Jones, E. E., & Goethals, G. R. (1972). Order effects in impression formation: Attribution context and the nature of the entity. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the cause of behavior* (pp. 27–46). Morristown, NJ: General Learning Press.
- Jones, E. E., & Nisbett, R. E. (1972). The actor and the observer: Divergent perceptions of the causes of behavior. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the cause of behavior* (pp. 79–94). Morristown, NJ: General Learning Press.
- *Jussim, L., Coleman, L. M., & Lerch, L. (1987). The nature of stereotypes: A comparison and integration of three theories. *Journal of Personality and Social Psychology*, 52, 536–546.
- Jussim, L., Nelson, T. E., Manis, M., & Soffin, S. (1995). Prejudice, stereotypes, and labeling effects: Sources of bias in person perception. *Journal of Personality and Social Psychology*, 68, 228–246.
- *Kameda, T. (1985). Stereotype-based expectancy and academic evaluation: The joint influence of prior expectancy and the diagnosticity of current information. *Japanese Psychological Research*, 27, 163–172.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction–integration model. *Psychological Review*, 95, 163–182.
- Klein, K., & Creech, B. (1982). Race, rape, and bias: Distortion of prior odds and meaning changes. *Basic and Applied Social Psychology*, 3, 21–33.
- Klein, W. M., & Kunda, Z. (1992). Motivated person perception: Constructing justifications for desired beliefs. *Journal of Experimental Social Psychology*, 28, 145–168.
- *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.
- Kruglanski, A. W., & Freund, T. (1983). The freezing and unfreezing of lay-inferences: Effects on impression primacy, ethnic stereotyping, and numerical anchoring. *Journal of Experimental Social Psychology*, 19, 448–468.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.
- Kunda, Z., Miller, D. T., & Claire, T. (1990). Combining social concepts: The role of causal reasoning. *Cognitive Science*, 14, 551–577.
- Kunda, Z., & Oleson, K. C. (1995). Maintaining stereotypes in the face of disconfirmation: Constructing grounds for subtyping deviants. *Journal of Personality and Social Psychology*, 68, 565–579.
- *Kunda, Z., & Sherman-Williams, B. (1993). Stereotypes and the construal of individuating information. *Personality and Social Psychology Bulletin*, 19, 90–99.
- Kunda, Z., Sinclair, L., & Griffin, D. (1995). *Equal ratings but separate meanings: Stereotypes and the construal of traits*. Unpublished manuscript.
- Langer, E. J., & Abelson, R. P. (1974). A patient by any other name . . . : Clinician group difference in labeling bias. *Journal of Consulting and Clinical Psychology*, 42, 4–9.
- *Linville, P. W., & Jones, E. E. (1980). Polarized appraisals of out-group members. *Journal of Personality and Social Psychology*, 38, 689–703.
- *Locksley, A., Borgida, E., Brekke, N., & Hepburn, C. (1980). Sex stereotypes and social judgment. *Journal of Personality and Social Psychology*, 39, 821–831.

- Locksley, A., Hepburn, C., & Ortiz, V. (1982). Social stereotypes and judgments of individuals. *Journal of Experimental Social Psychology*, 18, 23-42.
- Macrae, C. N., & Shepherd, J. W. (1989). Stereotypes and social judgments. *British Journal of Social Psychology*, 28, 319-325.
- Macrae, C. N., Sheperd, 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.
- Marr, D., & Poggio, T. (1976). Cooperative computation of stereo disparity. *Science*, 194, 283-287.
- McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review*, 88, 375-407.
- *McKirnan, D. J., Smith, C. E., & Hamayan, E. V. (1983). A sociolinguistic approach to the belief-similarity model of racial attitudes. *Journal of Experimental Social Psychology*, 19, 434-447.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81-97.
- Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and the invariance in personality structure. *Psychological Review*, 102, 246-268.
- Nisbett, R. E., Zukier, H., & Lemley, R. E. (1981). The dilution effect: Nondiagnostic information weakens the implications of diagnostic information. *Cognitive Psychology*, 13, 248-277.
- Pendry, L. F., & Macrae, N. (1994). Stereotypes and mental life: The case of the motivated but thwarted tactician. *Journal of Experimental Social Psychology*, 30, 303-325.
- *Pratto, F., & Bargh, J. A. (1991). Stereotyping based on apparently individuating information: Trait and global components of sex stereotypes under attention overload. *Journal of Experimental Social Psychology*, 27, 26-47.
- Pryor, J. B., & Ostrom, T. M. (1981). The cognitive organization of social information: A converging-operations approach. *Journal of Personality and Social Psychology*, 41, 628-641.
- Quillian, M. (1968). Semantic memory. In M. Minsky (Eds.), *Semantic information processing* (pp. 227-270). Cambridge, MA: MIT Press.
- Rasinski, K. A., Crocker, J., & Hastie, R. (1985). Sex stereotypes and social judgments: An analysis of the social perceiver's use of subjective probabilities. *Journal of Personality and Social Psychology*, 49, 317-326.
- Read, S. J., & Marcus-Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology*, 65, 429-447.
- Read, S. J., & Miller, L. C. (1993). Rapist or "regular guy": Explanatory coherence in the construction of mental models of others. *Personality and Social Psychology Bulletin*, 19, 526-541.
- Read, S. J., & Miller, L. C. (1994). Dissonance and balance in belief systems: The promise of parallel constraint satisfaction processes and connectionist modeling approaches. In R. C. Schank & E. J. Langer (Eds.), *Belief, reasoning, and decision making: Psychologic in honor of Bob Abelson* (pp. 209-235). Hillsdale, NJ: Erlbaum.
- Rosen, B., & Jerdee, T. H. (1974a). Effects of applicant's sex and difficulty of job on evaluations of candidates for managerial positions. *Journal of Applied Psychology*, 59, 511-512.
- Rosen, B., & Jerdee, T. H. (1974b). Influence of sex role stereotypes on personnel decisions. *Journal of Applied Psychology*, 59, 9-14.
- Rosenthal, R. (1991). *Meta-analytic procedures for social research*. Newbury Park, CA: Sage.
- 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.
- Rumelhart, D. E., & McClelland, J. L. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge MA: MIT Press/Bradford Books.
- 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.
- Secord, P. F., Bevan, W., & Katz, B. (1956). The Negro stereotype and perceptual accentuation. *Journal of Abnormal Social Psychology*, 53, 78-83.
- Shultz, T. R., & Lepper, M. R. (1992). A constraint satisfaction model of cognitive dissonance phenomena. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 462-467). Hillsdale, NJ: Erlbaum.
- Smith, E. R., & Miller, F. D. (1983). Mediation among attributional inferences and comprehension processes: Initial findings and a general method. *Journal of Personality and Social Psychology*, 44, 492-505.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99, 3-21.
- Spellman, B. A., & Holyoak, K. J. (1992). If Saddam is Hitler, then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, 62, 913-933.
- Stangor, C., Lynch, L., Duan, C., & Glass, B. (1992). Categorization of individuals on the basis of multiple social features. *Journal of Personality and Social Psychology*, 62, 207-218.
- Taylor, D. M., & Jaggi, V. (1974). Ethnocentrism and causal attribution in a South Indian context. *Journal of Cross-Cultural Psychology*, 5, 162-171.
- Taylor, S. E. (1981). A categorization approach to stereotyping. In D. L. Hamilton (Eds.), *Cognitive processes in stereotyping and intergroup behavior* (pp. 88-114). Hillsdale, NJ: Erlbaum.
- Taylor, S. E., Fiske, S. T., Etcoff, N. L., & Ruderman, A. J. (1978). Categorical and contextual bases of person memory and stereotyping. *Journal of Personality and Social Psychology*, 36, 778-793.
- Thagard, P. (1989). Explanatory coherence. *The Behavioral and Brain Sciences*, 12, 435-467.
- Thagard, P. (1992). *Conceptual revolutions*. Princeton: Princeton University Press.
- Thagard, P., & Millgram, E. (1995). Inference to the best plan: A coherence theory of decision. In A. Ram & D. B. Leake (Eds.), *Goal-directed learning* (pp. 439-454). Cambridge, MA: MIT Press.
- Thagard, P., & Verbeurgt, K. (1995). *Coherence*. Unpublished manuscript, University of Waterloo.
- *Triandis, H. C., Loh, W. D., & Levin, L. A. (1966). Race, status, quality of spoken English, and opinions about civil rights as determinants of interpersonal attitudes. *Journal of Personality and Social Psychology*, 3, 468-472.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review*, 93, 239-257.
- Trope, Y. (in press). Dispositional bias in person perception: A hypothesis-testing perspective. In J. Cooper & J. M. Darley (Eds.), *Attribution processes, person perception, and social interaction: The legacy of Ned Jones*. Hillsdale, NJ: Erlbaum.
- *Ugwuegbu, D. C. E. (1979). Racial and evidential factors in juror attribution of legal responsibility. *Journal of Experimental Social Psychology*, 15, 133-146.
- Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology*, 45, 961-977.
- *Whitehead, G. I., Smith, S. H., & Eichorn, J. A. (1982). The effects of

- subject's race and other's race on judgments of causality for success and failure. *Journal of Personality*, 50, 193-202.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneity of trait inferences. *Journal of Personality and Social Psychology*, 47, 237-252.
- Winter, L., Uleman, J. S., & Cunniff, C. (1985). How automatic are social judgments? *Journal of Personality and Social Psychology*, 49, 904-917.
- Wojciszke, B. (1994). Multiple meanings of behavior: Constructing actions in terms of competence or morality. *Journal of Personality and Social Psychology*, 67, 222-232.
- Wong, P., & Weiner, B. (1981). When people ask "why" questions, and the heuristics of attributional search. *Journal of Personality and Social Psychology*, 40, 650-663.
- Yarkin, K. L., Town, J. P., & Wallston, B. S. (1982). Blacks and women must try harder: Stimulus persons' race and sex attributions of causality. *Personality and Social Psychology Bulletin*, 8, 21-24.
- Yzerbyt, V. Y., Schadron, G., Leyens, J., & Rocher, S. (1994). Social judgeability: The impact of meta-informational cues on the use of stereotypes. *Journal of Personality and Social Psychology*, 66, 48-55.
- Zanna, M. P., & Hamilton, D. L. (1977). Further evidence for meaning change in impression formation. *Journal of Experimental Social Psychology*, 13, 224-238.

Appendix

IMPRESSION FORMATION (IMP)

IMP (short for "IMPRESSION formation") is a LISP program that we have used to simulate many important experimental findings concerning impression formation. The input to IMP is simple: (OBSERVED 'PERSON 'CHARACTERISTIC) is a LISP function that sets up a symmetric excitatory link between a special OBSERVED unit and a unit representing the CHARACTERISTIC. (ASSOCIATE 'CHAR1 'CHAR2 DEGREE) is a function that sets up a symmetric link between the two units representing the two characteristics CHAR1 and CHAR2. If DEGREE is positive, the resulting link is excitatory. If DEGREE is negative, the resulting link is negative. The default value for DEGREE is +1. In all the runs reported in this article, excitatory links have the default value of .04, and inhibitory links have the value of -.06.

To compute activations of the units, each unit is given a starting activation close to 0 (except the special OBSERVED unit, which is fixed at 1) and repeated cycles of updating begin. Activation is allowed to range between +1 and -1. On each cycle the activation of a unit j , a_j , is updated according to the following equation:

$$a_j(t+1) = a_j(t)(1-d) + \begin{cases} \text{net}_j(\max - a_j(t)) & \text{if } \text{net}_j > 0 \\ \text{net}_j(a_j(t) - \min) & \text{otherwise} \end{cases}$$

Here, d is a decay parameter (default = .05) that decrements each unit at every cycle, \min is a minimum activation (-1), \max is maximum activation (+1). Based on the weight w_{ij} between each unit i and j , one can calculate net_j , the net input to a unit, by:

$$\text{net}_j = \sum_i w_{ij} a_i(t).$$

When all units have reached stable activation values, which typically

happens after about 70 cycles (taking only a few seconds on a Macintosh Quadra 900), we view the activation level of a characteristic as indicative of the extent to which it characterizes the person. The output of IMP is thus activation values for units that represent the impression formed of a target person.

Note that IMP uses local connectionist representations in which each trait or stereotype is represented by a single unit. Biologically, it would be more plausible to use distributed representations in which a number of units correspond to a characteristic. We have not used distributed representations because no corpus of examples is available that could be used to train up a network with hidden units to provide a distributed representation of stereotypes. To create instances of members of a given stereotyped group, for example, we would have to use our assumptions about what characteristics these group members are likely to have. Therefore, the result of such training should be just to produce the positive and negative associations between characteristics that we encode directly. Moreover, the mechanisms by which stereotypes are learned may be different from the supervised learning algorithms usually used to produce distributed representations in PDP models.

Parallel constraint satisfaction can also be accomplished by using a nonlinear optimization algorithm that is guaranteed to satisfy at least 87.8% of the maximally satisfiable constraints (Thagard & Verbeugt, 1995). This algorithm, which uses a newly developed technique based on semidefinite matrix manipulation, is less biologically and psychologically plausible than connectionist algorithms that we use.

Received April 25, 1995

Revision received August 16, 1995

Accepted September 22, 1995 ■