Title

The dynamic relationship between stereotype efficiency and mental representation

Author

Sherman, J

Publication Date

2023-12-10

Peer reviewed

11

The Dynamic Relationship Between Stereotype Efficiency and Mental Representation

Jeffrey W. Sherman

Northwestern University

Stereotypes are "cognitive structures that contain the perceiver's knowledge, beliefs, and expectations about a human group" (Hamilton & Trolier, 1986, p. 133). In the past 20 years, a great deal of attention has been devoted to understanding how stereotypes influence the perception of individual members of stereotyped groups. This research has shown stereotypes to have a profound impact on many basic cognitive processes that underlie social perception. They determine the kinds of information to which people attend, the ways in which that information is interpreted and encoded, and the manner in which the information is stored in memory and subsequently retrieved. Of course, stereotypes also influence the content of resulting judgments about target individuals as well as perceivers' behavior toward those individuals (for a review, see Hamilton & Sherman, 1994).

More recently, attention has also turned to defining more precisely the nature of these stereotypic cognitive structures or mental representations (for reviews, see Hamilton & Sherman, 1994; Sherman, 1996). Broadly speaking, a *mental representation* is "an encoding of some information, which an individual can construct, retain in memory, access, and use in various ways" (Smith, 1998, p. 391). Research on the mental representation of stereotypes has been concerned with identifying the manner in which knowledge about social groups is constructed and retained in memory. When we say that a stereotype has been activated, what, specifically, do we mean? What exactly is being represented and activated in memory?

To this point, research on the influence of stereotypes on social perception processes and on the mental representation of stereotypes has proceeded in a largely independent fashion, with little overlap having been noted or pursued between the two topics. However, these topics are much more closely related than has been previously acknowledged. We cannot directly observe people's mental representations. We can only infer these representations through the processes that act on them. As such, any model of representation is also a model of cognitive processing and is constrained by what we know about cognitive processes. In turn, what we learn about mental representation constrains the viability of different process models. Thus, one cannot be fully understood without a corresponding understanding of the other (Anderson, 1978; Barsalou, 1990; Smith, 1998). The primary goal of this chapter is to detail how research on stereotype representation and stereotypic effects on information processing (particularly as they relate to stereotype efficiency) inform one another.

First, I describe recent research on stereotype representation and detail how concerns for efficient information processing play a vital role in the nature and development of stereotypic representations. Next, I explain how these representations in turn influence the processing of stereotype-relevant information. This discussion focuses on a recently developed model of how stereotypes confer efficiency in social perception. I argue that the efficient processing of stereotype-relevant information is largely dependent on the representational nature of stereotypes. Finally, I describe how these efficient processing mechanisms feed back into and reinforce the representational nature of stereotypes. Thus, the relationship between mental representation and processing efficiency is a dynamic one. Concerns about efficiency influence representation, which in turn influences processing efficiency, which in turn influences representation.

THE MENTAL REPRESENTATION OF STEREOTYPES

At one level, the mental representation of a stereotype can be defined in terms of the particular content that is included in the representation (e.g., the particular traits thought to characterize the target group). However, such a definition is merely descriptive and uninformative as to

the mechanisms by which stereotypes produce their effects. At another level, stereotype representation can be defined by the specificity with which the particular stereotypic traits are represented in memory. Research on this question has focused on the distinction between abstraction-based and exemplar-based knowledge (see Hamilton & Sherman, 1994; Sherman, 1996). Abstraction-based knowledge summarizes the features of a concept that have been extracted from experiences (both first- and second-hand) with specific instances of the concept. Generally speaking, these abstractions can be thought of as category averages along different attributes (including, perhaps, summary estimates of category variability; Hamilton & Sherman, 1994; Park & Hastie, 1987). In the case of knowledge about a social group, an abstract stereotype would consist of a group summary along a particular dimension (e.g., friendliness) that had been extracted from relevant experiences with the group and its members and had become associated with the group. In contrast, exemplar-based knowledge consists of the specific instances of the concept in question. Exemplar-based stereotypes would consist of representations of particular group members and their behaviors (e.g., specific friendly and unfriendly behaviors). A central argument of this chapter is that specifying the different contributions of these two types of knowledge to stereotype representation has significant implications for understanding the functional significance of stereotypes in information processing, and that defining stereotypes at this level of specificity can clarify the roles that stereotypes play in attentional, encoding, retrieval, and judgment processes.[1]

Historically, abstractionist views have held sway, with stereotypes being conceived of as (over)simplified generalizations about groups of people (e.g., Brigham, 1971; Lippmann, 1922). In more recent years, these generalizations have been described alternatively as *schemata* (e.g., Taylor & Crocker, 1981), *prototypes* (e.g., Brewer, 1988), *expectancies* (e.g., Hamilton, Sherman, & Ruvolo, 1990), and *Bayesian base rates* (e.g., McCauley & Stitt, 1978). Although there are subtle differences among these models, they share a most important feature in common: They all suggest that people develop and store abstract impressions that summarize the behavioral tendencies of social groups. These abstractions may initially be extracted from encounters with specific group members. However, once formed, they may be retrieved and used independently of information about particular individuals and their behaviors.

Although these conceptualizations are clearly abstractionist in nature, they are not highly specified models of mental representation. Specifying and testing underlying representational assumptions have never been primary concerns to researchers in these traditions. Indeed, few attempts have been made to verify the proposed abstract nature of these stereotypes. Moreover, the potential role of specific exemplars in group knowledge is usually not considered in these models. This is not to say that these models posit no role for exemplars. Rather, the question is simply never raised. Although these models have frequently been characterized as "pure-abstraction" models that are hostile to exemplar-based processes (e.g., Hamilton & Sherman, 1994; Sherman, 1996; Smith, 1990), they are so more out of neglect than out of clear intent.

Serious consideration of the role of exemplar-based knowledge in stereotype representation has appeared only recently. Researchers have proposed pure exemplar-based models of stereotyping that directly challenge traditional abstractionist conceptions (e.g., Linville, Fischer, & Salovey, 1989; Smith, 1990; Smith & Zarate, 1992). According to these models, stereotypes do not exist as independently stored knowledge structures. Rather, group knowledge consists solely of information about particular group member behaviors. If knowledge of the group as a whole is required, it must be created on the spot by retrieving and summarizing information about specific individuals.

A Developmental Mixed-Model of Group Representation

The preceding discussion of abstraction and exemplar-based models of stereotypes suggests that group knowledge is either entirely abstraction- or exemplar-based. However, it has become clear that neither pure abstraction nor pure exemplar models accurately describe group knowledge. It has been relatively straightforward to demonstrate that particular category exemplars may sometimes form the basis of people's perceptions of social groups (e.g., Sherman, 1996; Smith & Zarate, 1990) or at least influence those perceptions (e.g., Bodenhausen, Schwarz, Bless, & Wänke, 1995; Lewicki, 1985; Schwarz & Bless, 1992).[2] Likewise, it is clear that group knowledge is not always based on information about particular group members and is sometimes abstracted (Park & Hastie, 1987; Sherman, 1996; Sherman, Klein, Laskey, & Wyer, 1998). As a result, many theorists now subscribe to mixed models of mental representation. According to these models, knowledge of a group may be based on either abstract or exemplar information depending on the circumstances. For example, the extent to which group knowledge is based on

abstractions and exemplars has been posited to be mediated by group membership (ingroups vs. outgroups; Park & Judd, 1990; Park, Judd, & Ryan, 1991), group cohesion (Hamilton & Sherman, 1996), target stereotypicality (Sherman, Klein, et al., 1998), and learning order of abstract and exemplar information (Smith & Zarate, 1990), among other factors.

In my own research, I have proposed and tested a model of group representation that specifies group familiarity as a key determinant of the extent to which group knowledge is based on abstractions or exemplars (Sherman, 1996). According to this model, in the early stages of learning about a social group, judgments of the group are based on information about particular group members because too few exemplars have been encountered to support the formation of abstract knowledge. With sufficient experience with group members (or secondhand accounts of their attributes), perceivers form abstract representations of the attributes that are stereotypical of the group. Once formed, these abstractions may be retrieved independently from group exemplars to make judgments about relevant features of the group. Thus, group knowledge is exemplar based primarily when no abstract knowledge exists (although other factors may encourage exemplar based judgments once abstractions have developed; see also Klein & Loftus, 1993; Klein, Loftus, Trafton, & Fuhrman, 1992; Sherman & Klein, 1994, for familiarity-based models of the mental representation of self and other individuals).

The predictions of this model were supported in two experiments. In the first experiment, participants formed impressions of an ill-defined group of college students who all belonged to the same club on campus. No other information was given about the group that would allow participants to rely on preexisting group stereotypes in forming their impressions. After learning either a relatively small or large amount of information about the group, participants performed a judgment priming task that assessed the extent to which judgments about the group depended on the retrieval of information about particular group members. Briefly, this task required participants to either make a judgment about the group or perform a control task, and then retrieve a particular behavior performed by a member of the group. If participants spontaneously activate behaviors about specific group members when they make judgments about the group, then it should take relatively little time to retrieve one of these behaviors following the judgment task as compared with the control task. In this case, the behaviors would have already been activated and would be relatively accessible. The results demonstrate that group judgments were based on activated exemplars only when participants had received little information about the group. At high levels of group experience, judgments no longer depended on exemplar retrieval. In this case, it appeared that participants were able to rely on abstract, trait-based stereotypes that had been formed about the group.

The second experiment examined the mental representations of more familiar social groups. Participants learned either a relatively small or large amount of information about either a group of engineers or a group of priests. For these groups, participants may rely on preexisting stereotypes to form judgments of the groups along relevant dimensions, regardless of how little information they have received about particular group members. If these stereotypes exist as prestored abstractions, then judgments along relevant dimensions would not involve the retrieval of information about particular group members. However, if stereotypes are merely exemplar-based constructions, then relying on them would involve the retrieval of the novel, experimentally presented group members (e.g., Smith, 1990). That is, the most recently encountered exemplars would be activated in making the group judgment. The results demonstrate that judgments along stereotypical dimensions (engineer-intelligent, priest-kind) did not involve the retrieval of group exemplars, regardless of how little information had been presented about the groups. Even when participants had learned a small amount of information about the groups, judgments along these dimensions did not depend on exemplar retrieval, presumably because participants were instead relying on their preexisting abstract stereotypes. By contrast, judgments along nonstereotypical dimensions (engineer-kind, priest-intelligent) did rely on exemplar retrieval. This research suggests that stereotypes are often stored as abstract knowledge structures in memory, and that their availability reduces the role of exemplar-based processes (see also Sherman, Klein, et al., 1998).

The Need for Efficiency Influences Stereotype Representation

The developmental mixed model of stereotype representation is based on a functional analysis of memory. According to this analysis, the ability to summarize experiences across time, situations, and individuals is critical to the efficient operation of memory. There are two aspects to this efficiency. First, by capturing patterns of invariances in the environment, people are able to learn generalizable skills, predict the future, and explain novel events (McClelland, McNaughton, & O'Reilly, 1995; Nosofsky, Palmeri, & McKinley, 1994; Schank,

1982; Sherry & Schacter, 1987). In this way, information that is learned from past experiences can be brought to bear on a wide variety of novel experiences. Thus, an individual's behavior may be predicted and understood in light of the trait generalizations that have been made about the person. Likewise, a group member's behavior may be understood (or misunderstood) in light of the abstract stereotypes that have been formed about that person's group. This is what Bruner (1957a) referred to as *going beyond* the information given, and it is an important factor in the development of abstract knowledge.

The second reason that the development of abstractions is efficient has to do with the need for streamlined representations and cognitive processes. Theoretically, it would certainly be possible to create generalizations solely via repeated exemplar activation and summation. However, there are at least three reasons to posit that this would not be the case. First, the ability to create and apply knowledge that may be used in response to general and invariant features of the environment may be inhibited by the levels of temporal, spatial, and contextual detail preserved in exemplars (particularly behavioral episodes; Sherry & Schacter, 1987). As such, it would be useful to formulate and store abstract knowledge in addition to exemplars. Second, the predictive validity of these summarizations depends on the number of experiences included in the summarization. This is simply a question of sample size. Because each experience represents a single sample from the environment population, it is necessary to aggregate across many experiences to obtain a reliable population estimate (McClelland et al., 1995). This is related to the third point – that it is simply more efficient to maintain prestored abstractions than it is to re-create them every time they are needed by retrieving and summarizing specific events or exemplars. Most researchers assume that there are capacity constraints on the number of relevant events that may be found and summarized in a timely manner (Bruner, Goodnow, & Austin, 1956; Hamilton & Mackie, 1990; McClelland et al., 1995; Nosofsky et al., 1994; Rothbart, Fulero, Jensen, Howard, & Birrell, 1978). To constantly do so would be inefficient compared with keeping a running summary of relevant exemplars that could be accessed whenever necessary. Indeed, the retrieval and application of specific episodes is more easily disrupted than the application of abstract knowledge structures (e.g., Rothbart, et al., 1978; Sherman & Bessenoff, 1999, as described later; see Johnson et al., 1993; Tulving, 1983, for reviews). Thus, although it would be theoretically possible to retrieve and summarize large numbers of exemplars to generate valid population estimates, it would be inefficient to do so.[3]

This discussion makes clear the intimate relationship between form and function. The development, storage, and use of abstract knowledge are based, in part, on the need for efficient use of memory. Abstractions provide relatively large amounts of generalizable knowledge at relatively little cognitive expense, providing a healthy ratio of information gained to effort expended. Supporting this view is the impressive spontaneity with which perceivers draw abstractions from particular episodes. Past research has shown that perceivers derive abstract trait representations about the self and others relatively quickly and without being encouraged (e.g., Klein & Loftus, 1993; Klein et al., 1992; Sherman & Klein, 1994). Other research has shown that the leap from behavioral perception to trait inference occurs spontaneously, if not automatically (see Uleman, Newman, & Moskowitz, 1996, for a review). Related findings in a variety of experiments led Hastie and Park (1986) to conclude that it was difficult to invent situations in which abstraction processes did not spontaneously occur (see also Anderson, 1989).

Of most direct relevance to the concerns of this chapter, my work on stereotype representation (Sherman, 1996) demonstrated that much the same thing happens when perceivers are learning about social groups. As information about the group is acquired, trait summaries are extracted. Moreover, this research showed that group judgments were based on exemplars only when an abstract stereotype was not available. This suggests that there is a functional basis to the development of abstract stereotypes. It is more efficient to store and retrieve abstract stereotypes than it is to re-create group stereotypes anew each time they are needed by retrieving and summarizing information about particular group exemplars.

Comparisons to Other Models

Of course, the idea that stereotypes promote efficient social perception has a long history in social psychology. Both Lippmann (1922) and Allport (1954) focused extensively on the efficiency of stereotyping in their early analyses. More recent theoretical statements by Brewer (1988) and Fiske and Neuberg (1990) were also largely based on a functional analysis. Those models are specifically concerned with the extent to which judgments of individual group members are based on stereotypes about the groups to which they belong versus specific behaviors performed by the individuals. Both models rely on the distinction between "top-down" and "bottom-up" processing in their analyses. Whereas

top-down processing relies on preexisting knowledge (such as stereotypes) to deduce the properties of a new stimulus, bottom-up processing relies on raw perceptual input (such as observed behaviors) to induce the properties of a stimulus. Both models argue that top-down processing is more efficient than bottom-up processing in forming impressions of individuals.

Clearly there are important similarities between these models and the developmental/representational model described earlier (Sherman, 1996). Most significant, all three models argue that stereotype-based social perception is more efficient than perception based on more specific kinds of information (exemplars, behaviors).[4] However, there are important differences between the models as well. One obvious difference is that, whereas the developmental/representational model describes the formation of group impressions or stereotypes (although it may be generalized to impression of individuals; Klein et al., 1992; Sherman & Klein, 1994), the dual process models developed by Brewer (1988) and Fiske and Neuberg (1990) are concerned with the formation of impressions of individuals. A second difference is that, whereas the developmental/representational model focuses both on the roles of efficiency in the initial extraction of abstract stereotypes and in their subsequent use, the dual process models focus only on the efficiency of using stereotypes that are already fully developed. Little attention is paid in the dual process models to exactly how and why stereotypes develop the way they do. Finally, and most important, the developmental analysis in our model is particularly concerned with detailing the evolving representational basis of stereotypes and how it is related to stereotype efficiency. In contrast, although they are concerned with the relative contributions of top-down (stereotype-based) and bottom-up (behavior-based) processes in social judgments, the dual process models are less concerned with the underlying representation of the top-down knowledge that is being applied and the implications that it has for stereotype use. Fiske and Neuberg suggested that stereotypes may be either abstraction or exemplar based (e.g., Fiske, Neuberg, Beattie, & Milberg, 1987), but no discussion was devoted to identifying the conditions that determine the type of representation or what difference the two types of representation would make in stereotyping processes. Although Brewer (1988) clearly defined stereotypes as abstract group prototypes, little attention was given to how the abstract nature of stereotypes particularly influenced information processing. Thus, whereas the developmental/representational model is a model of both representation and judgment processes, the dual process models are primarily models of judgment processes. As we see later, the focus on underlying representation in the developmental/representational model provides insights into questions surrounding stereotype efficiency and encoding processes that are not addressed by the dual process models. First, however, a broader discussion of stereotype efficiency and encoding processes is warranted.

STEREOTYPE EFFICIENCY AND THE ENCODING OF SOCIAL INFORMATION

As described previously, it is widely accepted that stereotyping is more efficient than individuation (Bodenhausen et al., 1999; Brewer, 1988; Fiske & Neuberg, 1990; Hamilton & Sherman, 1994). This view has been supported by a considerable amount of research in recent years that has shown that stereotyping is particularly prevalent when attentional capacity is depleted. Whether due to physical depletion (e.g., Bodenhausen, 1990; Kim & Baron, 1988), task difficulty (e.g., Bodenhausen & Lichtenstein, 1987; Kruglanski & Freund, 1983; Pratto & Bargh, 1991), multiple task demands (Gilbert & Hixon, 1991; Macrae, Hewstone, & Griffiths, 1993), anxiety-induced arousal (e.g., Wilder & Shapiro, 1989), or positive moods (e.g., Stroessner & Mackie, 1993), situations that decrease the availability of processing resources have been shown to increase the extent to which perceivers rely on their stereotypes.

Miser Models: Stereotypes as Crutches

Historically, two different types of models have been proposed to account for the relative efficiency of stereotyping compared with individuation. First, many researchers have interpreted such findings from a *cognitive miser* perspective (e.g., Bodenhausen & Wyer, 1985; Brewer, 1988; Fiske & Neuberg, 1990; for an overview, see Fiske & Taylor, 1984). According to this orientation, perceivers are miserly with their resources and devote as little processing capacity to a given task as possible – what Allport (1954) called "the principle of least effort" (p. 173). The application of this analysis to stereotyping is further rooted in the heuristic processing tradition in the judgment and decision making literature, which suggests that people often rely on short-cuts or heuristics to achieve resource conservation (March & Simon, 1958; Tversky & Kahneman, 1974). As such, miser models of stereotyping suggest that perceivers rely on stereotypes as judgmental heuristics to relieve themselves of the cognitive effort of having to systematically process individuating information. Although these models are clear that perceivers may be more or less motivated to

process systematically (Brewer, 1988; Fiske & Neuberg, 1990), stereotype use is always equated with processing in a miserly and unmotivated fashion. Thus, as perceivers become more miserly, stereotype use increases and individuation decreases; as perceivers become more motivated, stereotype use decreases and individuation increases. Of course, as capacity-saving devices, stereotypes are also presumed to be particularly useful when resources are depleted and careful processing is difficult to achieve.

Given these models' historical basis in the judgmental heuristics tradition, it is not surprising that their primary focus has been on predicting the outcomes of judgment processes. In particular, they have been largely concerned with predicting the relative weights that will be given to category-based (i.e., stereotypes) and individuating information in determining the content of evaluative and descriptive judgments about group members. However, stereotypes and individuating information do not merely exert independent and separate influences on social perception. Rather, stereotypes and individuating information mutually influence and constrain one another. Of particular relevance to this chapter, stereotypes have been shown to have a significant impact on the manner in which stereotype-relevant individuating information is attended to, encoded, and retrieved. Furthermore, concerns for processing efficiency play an important role in determining exactly *how* stereotypes influence these processes. Though the miser models have not been so concerned with delineating these effects, other models have made them of central importance.

Stereotypes as Filters

The second type of model that has been proposed to explain stereotype efficiency is the filter model (e.g., Bodenhausen, 1988; Bodenhausen & Lichtenstein, 1987; Bodenhausen, Macrae, & Garst, 1997; Hamilton & Sherman, 1994; Macrae et al., 1993; Macrae, Milne, & Bodenhausen, 1994; Macrae, Stangor, & Milne, 1994; Miller & Turnbull, 1986; Stangor & Duan, 1991; Stangor & McMillan, 1992; Taylor & Crocker, 1981). In contrast to miser models, filter models are largely concerned with how stereotypes influence the processing of individuating information. Also, in contrast to miser models, filter models do not suggest that increased stereotype use is always accompanied by diminished reliance on individuating information. Rather, certain kinds of individuating information benefit from the use of stereotypes. Based on schematic principles of memory (e.g., Minsky, 1975; Neisser, 1976), these models suggest that information that is consistent with a stereotypic expectancy is encoded and represented in memory more completely than information that is unrelated to the expectancy. In contrast, information that is inconsistent with the expectancy is particularly unlikely to be successfully encoded into memory. Thus, stereotypes act as filters that let in information that confirms people's expectancies and keep out information that disconfirms them.

There are two varieties of filter models. The weak or passive filter model suggests that the encoding advantage for consistent information is based on the fact that, because it fits with existing expectancies, it is simply easier to comprehend than inconsistent information. By providing this conceptual fluency to consistent information, stereotypes reduce the amount of capacity necessary to encode that information, freeing up processing resources for other tasks.

The strong or active filter model further suggests that, because consistent information is easier to encode, attention is directed toward consistent and away from inconsistent information (Bodenhausen, 1988; Bodenhausen & Lichtenstein, 1987; Bodenhausen et al., 1997; Fiske & Neuberg, 1990; Hamilton & Sherman, 1994; Macrae, Milne, & Bodenhausen, 1994; Stangor & Duan, 1991; Taylor & Crocker, 1981). Thus, stereotypes further promote efficiency by directing resources away from information that is difficult to encode and toward information that is easy to encode. This attentional hypothesis is based on the cognitive miser analysis that people generally prefer to do as little cognitive work as necessary. It is also derived, in part, from principles of selective exposure (e.g., Festinger, 1957; Frey, 1986), which suggest that people prefer to not attend to information that challenges their beliefs (see also Thompson, Naccarato, Parker, & Moskowitz, chap. 2, this volume), particularly if they do not possess the resources to refute that information. Indeed, it is important to note that, as is the case with miser models, both filter models predict that the influence of stereotypes is especially likely to be observed when processing resources are low and efficient processing is at a premium. Thus, in these conditions, the encoding advantages for consistent over inconsistent information should be particularly strong. However, in contrast to miser models, these models suggest that perceivers who are particularly motivated to process target information may be especially likely to rely on their stereotypes to arrive at a preferred impression.

STEREOTYPE REPRESENTATION, ENCODING FLEXIBILITY, AND STEREOTYPE EFFICIENCY

Recently we have proposed and tested an alternative

model of stereotype efficiency that makes different predictions from both the miser and filter models (Sherman & Frost, 2000; Sherman, Lee, et al., 1998). This model is particularly concerned with the manner in which stereotype-relevant individuating information is attended to and encoded, and how concerns for efficiency influence these processes. The predictions for this model were derived in part from our earlier work on the mental representation of stereotypes (Sherman, 1996), in which the presence of an abstract stereotype was shown to reduce the extent to which stereotype-consistent group judgments were based on particular stereotype-consistent behaviors. In conditions in which perceivers possessed abstract stereotypes, they were freed from having to rely on specific stereotypical behaviors as the basis for their judgments. Note that the implications of these findings are opposed to the predictions of filter models (particularly strong filter models). Whereas those models suggest that stereotypes increase the encoding and influence of stereotype-consistent behaviors, our results suggest that stereotypes decrease the importance of carefully encoding stereotypical behaviors. Because stereotypical information is already provided by abstractions, the necessity of attending to, carefully encoding, and retrieving confirmatory behaviors is reduced (see also Johnston & Hawley, 1994; von Hippel, Jonides, Hilton, & Narayan, 1993; von Hippel, Sekaquaptewa, & Vargas, 1995). Moreover, whereas filter models suggest that encoding should be particularly biased toward expected information when capacity is depleted, our results imply that careful encoding of expected information is particularly unlikely under these conditions. If, as we have suggested, it is more efficient to rely on abstract knowledge than behavioral exemplars, then it would seem unlikely that stereotypical behaviors would be particularly well encoded when processing resources are scarce.

What Is an Efficient System?

However, our model is not a miser model. We do not wish to suggest that perceivers rely on stereotypes out of laziness or that stereotype use necessarily implies the disuse of individuating information. On the contrary, rather than minimizing effort, we argue that stereotype use is about maximizing efficiency – the amount of information gained for effort expended. Indeed, the other major impetus for our model was a theoretical consideration of what an efficient cognitive system ought to accomplish. A number of researchers have argued that cognitive systems that are either too stable or too flexible would be at an evolutionary disadvantage (e.g., Johnston & Hawley, 1994; Sherry & Schacter, 1987; Tulving, Markowitsch, Kapur, Habib, & Houle, 1994). On the one hand, if knowledge is too easily changed, then it is not useful for making consistent predictions about the environment. Hence, there must be some stability in our expectations. On the other hand, if expectations are not open to alteration, they may not adequately represent the world. If expectancies have no predictive value, then they also are not very useful. At some point, people must be able to recognize that the data do not fit their expectancies, and they must revise them. Thus, for maximum predictive value, efficient systems must encode both invariances in the environment, which encourage the development of expectancies, and variances (unexpected events), which suggest that expectancy reorientation may be necessary (e.g., Johnston & Hawley, 1994; McClelland et al., 1995; Nosofsky et al., 1994; Schank, 1982; Sherry & Schacter, 1987; Tulving et al., 1994). Given the importance of each of these goals, it might be expected that there would be provisions in the cognitive system for encoding both expected and unexpected information when processing resources are scarce.

This functional analysis is at odds with both miser and filter models. Both of those types of models propose cognitive systems that are inherently conservative. According to miser models, when resources are depleted, perceivers rely on their existing stereotypic expectancies to the relative exclusion of all novel behavioral information. Strong filter models suggest that, when capacity is low, people focus on confirmatory behaviors and direct their attention away from disconfirming information. Finally, weak filter models propose that, when resources are low, perceivers are unable to encode unexpected behaviors because they are simply too difficult to comprehend. Although people may want to encode this information, they are simply unable to do so. Thus, in all three models, the overwhelming trend is toward maintaining stability in existing expectations, particularly when capacity is restricted. In contrast, a model based on the notion of an efficiency-maximizing system suggests that efficient stereotype use ought to facilitate, in different ways, the encoding of both expected and unexpected individuating information when capacity is low.

Encoding Flexibility and Stereotype Efficiency

Recently, we have proposed and tested such an encoding flexibility model of stereotype efficiency (Sherman, Lee, et al., 1998). According to this model, stereotypes facilitate the encoding of consistent behaviors by providing explanatory frameworks that render those behaviors conceptually fluent. As such, those behaviors are relatively easy to comprehend, even when processing

resources are scant. This aspect of our model coincides with the passive filtering hypothesis. However, our model suggests that the conceptual fluency of consistent information has a different impact on the encoding of consistent and inconsistent information than do filter models. In particular, we argue that, because consistent information can be understood with relatively little effort, substantial attention is not devoted to encoding the details of this information (e.g., Graesser, 1981; Johnston & Hawley, 1994; Sherman, 1996; von Hippel et al., 1993), particularly when capacity is depleted. Instead, those resources may be redirected to assist in the encoding of inconsistent information, which is difficult to understand. Yet this does not mean that the inconsistent information is fully understood – only that the effort will be made. Thus, when capacity is low, conceptual encoding (encoding for gist meaning) favors consistent information, whereas attentional allocation and perceptual encoding (encoding for details) favor inconsistent information. As such, stereotypes do not merely simplify impression formation for lazy perceivers. Rather, they permit the flexible distribution of resources in a way that maximizes the amount of information gained for the effort expended. This encoding flexibility is functional because it promotes both stability and plasticity in the mental system.

The results from a number of experiments have supported the predictions of this model. Three experiments tested our hypotheses about the distribution of attentional resources under different encoding conditions (Sherman, Lee, et al., 1998). In all three experiments, participants read stereotype-consistent and -inconsistent information about a target person while either under a cognitive load or not. The first experiment measured the amount of time participants spent reading the consistent and inconsistent behaviors as a function of processing capacity. Results show that participants spent an equal amount of time reading consistent and inconsistent information when capacity was high, but spent a greater amount of time reading inconsistent than consistent information when capacity was low. Thus, participants were not directing their attention away from the inconsistent information when they were under load. The second experiment used a dual-task paradigm to examine the amount of attention paid to consistent and inconsistent information as a function of cognitive capacity. As participants read about the target, they were also asked to monitor auditory tones emitted by their computers. When capacity was high, participants responded to this secondary task equally quickly, regardless of whether it occurred during the encoding of consistent or inconsistent items. In contrast, when capacity was low, participants took more time to respond to the secondary task when it occurred as inconsistent items were being encoded than when consistent items were being encoded. This response interference demonstrates that greater attention was devoted to encoding the inconsistent than consistent items when capacity was depleted. Finally, the third experiment forced participants to attend selectively to either consistent or inconsistent information by presenting consistent and inconsistent items in pairs for a brief period of time. The results show that, when capacity was high, participants recognized consistent and inconsistent items from a pair equally well. However, when capacity was depleted, the inconsistent item in the pair was recognized with significantly greater accuracy than the consistent item. Together these three experiments provide strong evidence that, when processing capacity is limited, greater resources are devoted to the encoding of inconsistent than consistent information, which is in direct opposition to the predictions of filter models.

Two other experiments tested our hypotheses about the perceptual and conceptual encoding of stereotype-consistent and -inconsistent information (Sherman, Lee, et al., 1998). Once again, participants read stereotype-consistent and -inconsistent information about a target person while either under a cognitive load or not. Subsequently, they engaged in a priming task that measured either perceptual or conceptual encoding. Both tasks required participants to identify words that were flashed very briefly (33 ms) on computer screens. The perceptual priming task examined the extent to which participants could identify words that had appeared in the original stimulus items, but were unrelated to the gist meaning of the items (e.g., the word *salesgirl* from the sentence, "Swore at the salesgirl"). Participants' ability to identify these words reflects the extent to which the perceptual details of the items had been extracted during encoding. In contrast, the conceptual priming task examined the extent to which participants could identify words that reflected the gist meaning of the original stimulus items, but had not actually appeared in those items (e.g., trait terms such as *kind* and *mean*). Participants' ability to identify these words reflects the extent to which the gist meaning of the items has been extracted during encoding (for a methodological overview, see Roediger, 1990).

The first experiment showed that perceptual encoding was greater for inconsistent than consistent behaviors under conditions of both high- and low-processing capacity. The second experiment showed that, under high-capacity conditions, conceptual encoding was equally strong for consistent and inconsistent behaviors. In contrast, when encoding capacity was limited, the conceptual meanings of consistent behaviors were much

more likely to be extracted than the conceptual meanings of inconsistent behaviors. Thus, despite the attentional and perceptual encoding advantages for inconsistent information when resources were low, conceptual encoding favored consistent information under such conditions.

These five experiments provide strong initial support for the encoding flexibility model of stereotyping. When resources are limited, stereotypes facilitate the encoding of both stereotype-consistent and -inconsistent information. Inconsistent information receives greater attention and more thorough perceptual encoding. However, despite these advantages, conceptual encoding favors consistent information in these same conditions. Through these encoding flexibility processes, when resources are scarce, stereotypes are able to promote their own stability (through conceptual encoding) while maintaining vigilance (through attentional distribution and perceptual encoding) that reorientation may become necessary.

We argue that this encoding flexibility is facilitated by an abstract representational structure of stereotypes. As described previously, well-developed stereotypes are stored as abstract knowledge structures that summarize group-level information across time, situations, and individuals (Sherman, 1996). As such, they may aid in understanding novel behaviors performed by novel group members. This is particularly true for behaviors and individuals that are congruent with the basic gist of the stereotype. Because the information provided by these acts/individuals is redundant with, and may be assumed by, the broader abstraction, perceivers need not attend carefully to it. Instead, those resources may be directed to assist in the encoding of other information that does not fit so comfortably with the stereotype (e.g., unexpected information), particularly when resources are low.

If ever, it may be in the early stages of stereotype development, when group-based expectancies are too weak to support the formation of abstract stereotypes, that attentional allocation and perceptual encoding would be biased toward stereotype-confirming information (e.g., Klayman & Ha, 1987; Skov & Sherman, 1986). In such cases, the stereotype is less useful for inferring the meanings of novel stereotypical behaviors. As such, particular stereotype-consistent acts demand greater attention. Moreover, the stereotype may not present a strong enough expectancy to produce clearly identifiable, inconsistent data that would draw special attention (e.g., Hamilton & Sherman, 1996; Sherman, 1996; Sherman & Klein, 1994; Srull, Lichtenstein, & Rothbart, 1985). Finally, it may be that information that facilitates the establishment of an abstract expectancy that allows for future generalization may receive more careful processing than information that challenges an expectancy that is weak and ungeneralizable to begin with. Thus, prior to the completion of the abstraction process, consistent information may attract more attention and perceptual encoding than inconsistent information. However, as stereotypic expectancies congeal and become abstract, consistent information becomes redundant and may be readily inferred from the stereotype. At the same time, inconsistent information becomes easier to identify and gains in importance. As a result, resources are more likely to shift from consistent to inconsistent information.

EFFICIENT PROCESSING INFLUENCES STEREOTYPE REPRESENTATION

First I described how concerns for efficiency influence the initial development of abstract stereotypes. Next I showed that, once these abstractions are formed, they influence the efficiency with which stereotypes impact the subsequent encoding of stereotype-relevant information. Now I take the discussion full circle and describe how efficient encoding flexibility processes feed back into and influence the mental representation of stereotypes and stereotype-relevant information.

One of the more important questions arising from our research on encoding flexibility has to do with the mechanisms of stereotype plasticity. The attentional and perceptual encoding advantages for inconsistent information under cognitive load must ultimately contribute to stereotype revision, but how? Among other reasons, we believe these encoding advantages are important for stereotype plasticity because they help people retrieve and reconstruct the details of unexpected events at a later time when greater resources are available for comprehension. Because these items of information are difficult to interpret during encoding (especially if resources are low), it is of particular importance to be able to retrieve their details at a later time, when their implications may be more fully understood. Retaining the details of these behaviors would also be important so that they may be compared to newly observed behaviors that may help clarify matters and promote conceptual consolidation (e.g., McClelland et al., 1995). In fact, a number of researchers have argued that the essential purpose of episodic memory is to record the details of unexpected events for later inspection and comparison (McClelland et al., 1995; Nosofsky et al., 1994; Schank, 1982; Sherman & Bessenoff, 1999; Sherry & Schacter, 1987). In contrast, retaining the specific details of expected information is not such a pressing matter. In this case, the basic gist may simply be extracted and stored as semantic memory.

Our demonstrations of attentional and perceptual encoding advantages for stereotype-inconsistent information are certainly consistent with this functional analysis of episodic and semantic memory. However, we may also make specific predictions about the extent to which the encoding of stereotype-consistent and -inconsistent behaviors would encourage the development and use of abstraction-based versus exemplar-based group knowledge. In particular, if people simply extract the basic gist of consistent behaviors and store it as semantic memory, then we would expect perceivers to quickly develop abstract group knowledge pertaining to those behaviors. This is essentially what Sherman (1996) showed. In contrast, if people are more concerned with retaining the perceptual and contextual details of inconsistent behaviors, then we would expect that the development and use of abstract knowledge pertaining to those behaviors would be substantially delayed, and the role of exemplar-based processes would be enhanced.

The results of two recently reported experiments support these hypotheses (Sherman, Klein, et al., 1998). Participants were first randomly assigned to one of two social categories (overestimators or underestimators) based on a minimal group manipulation (Howard & Rothbart, 1980). Next, they were presented with positive or negative descriptors of members of either their ingroup or their outgroup. Finally, participants performed the previously described judgment priming task that assessed the extent to which group judgments depended on the retrieval of information about particular group members (Sherman, 1996). When the descriptors confirmed participants' expectancies that ingroups would be positive and outgroups would be negative, judgments about the groups did not involve the activation of specific exemplars. Instead, judgments were apparently based on abstract group summaries created during the encoding of the expected behaviors. In contrast, when the stimulus information suggested that the ingroup was negative or the outgroup was positive, judgments of the groups were constructed by retrieving from memory specific group exemplars (for related results, see Maass, Milesi, Zabbini, & Stahlberg, 1995; Maass, Salvi, Arcuri, & Semin, 1989). This implies that perceivers did not form semantic summaries during the encoding of unexpected behaviors. These data suggest that, if perceivers are unable (or unwilling) to extract the basic gist of unexpected behaviors, those episodes may be stored and retrieved for future use. Based on our other work on encoding flexibility, we might expect that these encoding/representational differences between consistent and inconsistent behaviors would be particularly evident when processing capacity is restricted during the initial encoding of group behaviors, and inconsistent information is particularly difficult to comprehend.

These studies demonstrate how the differential processing of stereotype-consistent and -inconsistent information may influence the manner in which that information gets represented in memory. Because perceivers are relatively unconcerned with encoding the specific details of consistent behaviors, that information may be "semanticized" and stored as abstract knowledge relatively quickly. By contrast, perceivers' concerns with encoding the details of inconsistent behaviors delay the abstraction process and enhance the role of episodic memory in the representation of such information. Thus, not only does the initial abstract nature of stereotypes influence the manner in which stereotype-relevant behaviors are encoded, but those encoding processes also subsequently feed back into and influence the mental representations of stereotypes and stereotype-relevant information. In both cases, concerns for efficiency are an important contributing factor.

ON THE NEED FOR ENCODING FLEXIBILITY

I have argued that it is important to encode and represent the perceptual and contextual details of stereotype-inconsistent but not -consistent information. In part, this is because the meanings of inconsistent behaviors are relatively difficult to discern during encoding. As such, it is useful to be able to reconstruct the details of those behaviors at a later time when further attempts at comprehension may proceed. However, there is another important reason that perceivers would more carefully encode the details of inconsistent than consistent behaviors. If the details of consistent behaviors are forgotten or are temporarily inaccessible, perceivers may simply rely on their abstract stereotypes to infer the basic gist of what has occurred. On the contrary, if the details of inconsistent behaviors cannot be retrieved, the meaning conveyed by those behaviors may also be lost if the initial conceptual encoding of those behaviors has been less than ideal. This is because the basic gist of inconsistent information cannot be reconstructed from existing abstract group knowledge after the fact. Thus, the careful encoding of the details of unexpected information not only helps advance the ultimate conceptual consolidation of that information, but also acts as insurance against the irreplaceable loss of such information. We conducted an experiment to examine this issue in more detail (Sherman & Bessenoff, 1999).

One important kind of contextual information is source information. To fully benefit from episodic

memory, perceivers must be able to accurately attribute events to their proper source (e.g., Johnson et al., 1993). Was the source of a news story the *New York Times* or the *World Weekly News*? Was that great research idea mine or my graduate student's? Was it John or Juan who threw the first punch in the bar brawl? In this last case, stereotypes about the aggressiveness of Hispanic men may influence the likelihood that people attribute the first punch to Juan rather than John. In this experiment, we examined such source attributions for stereotype-consistent and -inconsistent information. In particular, we were interested in the extent to which perceivers would misattribute stereotype-consistent and -inconsistent behaviors to a person who did not, in fact, perform them.

After reading lists of stereotype-consistent and -inconsistent behaviors that did and did not describe a target person, participants were asked to engage in a source monitoring task that required them to accurately attribute the behaviors to their proper source. Some participants were placed under a cognitive load as they performed this source monitoring task. Results show that, when the true target behaviors and false behaviors shared a similar encoding context (i.e., were encoded in the same time and place), and were therefore easily confused, participants relied on their stereotypes as cues in the source monitoring task and misattributed more false stereotype-consistent than -inconsistent behaviors to the target. Because the proper source of the behaviors was difficult to discern under these conditions, participants relied on their stereotypes as judgmental cues (see also Banaji & Greenwald, 1995). However, an interaction with the processing capacity variable demonstrated that this effect was only found when participants' ability to rely on episodic memory was impaired by the imposition of the cognitive load. Thus, when participants needed to rely on episodic memory because there was source/context confusion, but were unable to do so because of the imposed cognitive load, they relied on the stereotype as a heuristic cue in making their source attributions. In contrast, the stereotype was not used as a source cue (i.e., an equal number of misattributions were made for stereotype-consistent and -inconsistent behaviors) when participants possessed full processing capacity, although source confusion was high. In this case, participants relied on a more systematic analysis of episodic memory to reconstruct the source information about the behaviors and make their attributions.

These results highlight the necessity of carefully encoding and representing the episodic details of unexpected behaviors. When episodic memory is not available to perceivers, as in the cognitive load condition of this study, they may yet rely on abstract knowledge to guide social cognition. In such cases, the basic gist of stereotype-consistent behaviors may be inferred even if the details of those behaviors are inaccessible. In contrast, the ability to rely on abstract representations to fill in knowledge gaps does not extend to stereotype-inconsistent behaviors, for which there are no relevant abstractions. Therefore, it is of critical importance to retain the episodic details of these behaviors, particularly to the extent that their conceptual meanings have not been extracted during encoding. This suggests why the encoding flexibility processes outlined earlier (i.e., directing attentional and perceptual encoding toward unexpected behaviors when capacity is depleted) may be pursued in the first place: to ensure the adequate representation of unexpected events that may not be reconstructed if lost.

SUMMARY AND FUTURE DIRECTIONS

That mental representations and cognitive processes are mutually constraining in memory and judgment processes is axiomatic. Nevertheless, social-cognitive research on representation and processing has frequently proceeded in an independent fashion, with too few explicit attempts having been made to take stock of what the findings in one domain imply for theoretical development in the other domain (notable exceptions include Carlston, 1994; Klein & Loftus, 1993; Smith, 1990, 1996; Smith & DeCoster, 1998; Srull & Wyer, 1989; among others). The main goal of this chapter has been to demonstrate the utility of analyzing representation/processing dynamics in the context of stereotyping research.

I began by describing research aimed at specifying the representational basis of stereotypical knowledge. The initial theoretical impetus for this research and the representational model that evolved from it were based, in large part, on a functional analysis of efficient information processing, and how it may be facilitated by the development of generalizable knowledge structures (e.g., Bruner, 1957b; Klein & Loftus, 1993; Klein et al., 1992; Lippmann, 1922; McClelland et al., 1995; Nosofsky et al., 1994; Schank, 1982; Sherman & Klein, 1994; Sherry & Schacter, 1987). Thus, our theorizing about representation was heavily influenced by existing theories of processing from the outset.

Subsequently, I described our recently developed model of stereotype efficiency and encoding flexibility. The development of this model was heavily influenced by the conclusions drawn from our initial work on stereotype representation. The representational perspective provided by that initial research enabled us to make novel (and sometimes nonintuitive) predictions about

stereotype-based encoding processes that had not been identified by prior models of stereotype efficiency. Thus, in this case, it was the knowledge of representational issues that informed the research on processing.

Next, I presented research demonstrating how concerns about efficient processing and the differential encoding of stereotype-consistent and -inconsistent information feed back into and influence the mental representation of stereotype-relevant information. Thus, not only does the representational nature of stereotypes influence processing, but so too do processing strategies influence the representational nature of stereotypes. Finally, I described an experiment that further illustrated how the representational nature of stereotypes influences encoding processes and why it is that efficient encoding proceeds in the manner that it does. Altogether, the research described in this chapter makes a strong case for the value of integrating analyses of mental representation and information processing.

On Efficiency

One common factor running through all of the research described in this chapter is that stereotypes and stereotyping are influenced by perceivers' need for cognitive efficiency. Both the initial development and nature of stereotypes as well as their subsequent uses in social perception are molded by such concerns. In no way is this meant to suggest that other factors (e.g., identity-related motives, ego-defensive motives, etc.) are not critically important in stereotyping processes. Nor is this meant to imply that efficiency concerns are necessarily the most important factor in understanding stereotypes. Nevertheless, it is clear that efficiency is one critical factor that must be considered in models of stereotype representation and use (cf. Bodenhausen et al., 1999; Kunda & Thagard, 1996).

Although the important role of cognitive efficiency in stereotyping has long been recognized (e.g., Allport, 1954; Brewer, 1988; Fiske & Neuberg, 1990; Hamilton & Sherman, 1994; Lippmann, 1922), the findings outlined in this chapter suggest that some refinements may be needed in how stereotype efficiency is conceptualized. The efficient stereotype typically has been conceived of as either a crutch or a filtering device. Both metaphors stress the effort-reducing and simplifying properties of stereotypes as well as the conservative, stereotype-confirming nature of efficient processing. However, efficiency is defined here as the ratio of information gained (production) to effort expended (cost). As such, considerations of stereotype efficiency must consider not only savings in effort but also gains in information. Our research (Sherman & Frost, 2000; Sherman, Lee, et al., 1998) demonstrates that the primary motive behind stereotyping when capacity is depleted is not always simplicity for simplicity's sake. Perceivers are not so much interested in effort reduction per se as they are in efficiently distributing their resources. Furthermore, all aspects of efficient encoding are not motivated by, and do not necessarily result in, stereotype reconfirmation. Rather, both stereotype stability and plasticity are reinforced in different ways through the efficient use of stereotypes.

It is important to emphasize that, in contrast to claims made by some researchers (e.g., Oakes & Turner, 1990; Spears & Haslam, 1997), the idea that stereotypes are used to enhance rather than reduce information processing is in no way inconsistent with the view that stereotype use is often influenced by concerns for efficiency (for further discussion of this matter, see Sherman, Macrae, & Bodenhausen, 2000). To the contrary, stereotype contributions to information gain are one critical component of their efficiency, as demonstrated by our research. It may be true that researchers have tended to place greater emphasis on the effort reduction than the information gain side of the equation. However, our research in no way challenges (and, in fact, strongly supports) the basic point that stereotype use is influenced by the need for efficiency.

Implications for Dual-Process Models of Stereotyping. The model of efficiency we have outlined and the accompanying data have important implications for a number of central issues in stereotyping research. Dual process models have been developed to account for the conditions under which judgments about individual group members are dependent on top-down, stereotype-driven processes versus bottom-up integration of target behaviors (see Bodenhausen et al., 1999, for a review). In both of the prominent dual process models of stereotyping (Brewer, 1988; Fiske & Neuberg, 1990), increases in stereotype use are associated with decreases in individuation, particularly in the amount of attention paid to and the encoding quality of stereotype-inconsistent information. That is, because perceivers are using their stereotypes, they need not expend extra resources on difficult individuating processes. However, our results show that decreases in processing capacity increased both stereotyping processes (via conceptual encoding) and certain individuating processes (attentional allocation toward and perceptual encoding of inconsistent compared with consistent information) at the same time. These results demonstrate that efficiency-driven stereotype use need not exclude individuation processes. In fact, they suggest that stereo-

type use and individuation should be conceived as two separate but related continua, rather than as mutually exclusive processing modes (see Chaiken, Liberman, & Eagly, 1989). As such, movement along the two continua may proceed along different dimensions of encoding at the same time. Thus, stereotyping may be increased via one mode of encoding (e.g., conceptual), while individuation is increased via a different mode of encoding (e.g., perceptual) simultaneously.

The Interaction of Motivation and Processing Capacity in Stereotyping. In a related matter, our findings also shed light on the relationship between perceiver motivations and processing capacity in stereotyping. Dual process models of stereotyping suggest that motivations to individuate a target may override or diminish the influence of stereotypes (see Monteith, chap. 23, this volume). Yet these and other models of stereotype efficiency suggest that, when processing capacity is depleted, the impact of perceivers' processing goals are limited. The working model has been that motivations may be realized only if sufficient capacity is available. Thus, perceivers who may have initially been motivated by accuracy to individuate a target are presumed to abandon that motivation when resources are scant, and revert to more miserly pursuits such as purely heuristic processing or schematic filtering processes. In contrast, our results suggest that motivations continue to exert an important influence when cognitive resources are depleted (see also Moskowitz, Gollwitzer, Wasel, & Schaal, 1999). In our efficiency studies, participants were largely motivated by concerns for accuracy and did not forsake those concerns when capacity was diminished. Rather, they distributed their resources in such a way that enabled them to flexibly encode different aspects of both stereotype-consistent and -inconsistent information. Of course, there are many situations in which perceivers are motivated by concerns for ego-defense or social identity rather than by accuracy. In these contexts, we would expect that perceivers would, in fact, be especially likely to engage in miserly/filtering processes if resources are scant. Thus, stereotypes are flexible tools that are adapted to the current goals of the user. However, resource scarcity does not determine the processing goal, but rather how an already chosen goal is pursued.

Conclusion

The preceding analyses suggest that stereotypes are much more versatile tools than crutches or filters. Perhaps the metaphor of the Swiss Army knife, with multiple tools working simultaneously, better captures the flexibility with which efficient processing may be pursued. In terms of stereotyping, the tools correspond to different attentional, encoding (including perceptual and conceptual encoding), and inference processes that proceed in parallel. On some occasions, the tools may be working in concert toward achieving a single goal (e.g., stereotype confirmation). On other occasions, different tools may be used to pursue different goals at the same time (as in the case of encoding flexibility processes).

One of the major challenges for future research on stereotyping should be to more closely examine how different processing conditions (e.g., variations in attentional capacity) and different processing motives (e.g., accuracy vs. defense vs. self-presentation) interact to determine how perceivers use stereotypes in forming impressions of others. This challenge must be met with a more complete understanding of cognitive efficiency that takes into account both effort reduction and information gain strategies, as well as both stability-maintaining and plasticity-seeking processes. Finally, progress in these endeavors can be greatly advanced by further attempts to integrate what is known about the mental representation of stereotypes and the processing of stereotype-relevant information.

ACKNOWLEDGMENTS

Preparation of this article was supported by NIMH Grant 55037 to Jeffrey W. Sherman. Thanks to Galen Bodenhausen, Jim Sherman, and Eliot Smith for their thoughtful comments.

ENDNOTES

1. At still a third level, the mental representation of a stereotype may be defined in terms of a particular representational format (e.g., associative network models vs. localist connectionist models vs. distributive connectionist models). This level of analysis attempts to define at some hypothetical neural level the manner in which knowledge about a group is stored in memory. All of the prominent format models describe a memory system that is based on activation spreading between associated quasineuronal units or nodes that represent concepts or their features (e.g., connections or overlapping patterns of connections between social category representations and stereotypical features). Although the different models share much in common in terms of their basic metaphors, they make different assumptions about the nature of memory nodes, associations, and the ways that activation spreads through a representation. A full discussion of the implications of each of these models for stereotyping goes well beyond the scope of this chapter

(for an excellent review, see Smith, 1996). However, it is important to note that each type of format can (and must) account for the storage and retrieval of both exemplars and abstractions. As such, the issues surrounding the differential impact of abstraction- and exemplar-based stereotypes discussed in this chapter are largely independent of questions of format.

2. Note that demonstrations that accessible exemplars may influence group judgments do not indicate that abstract group impressions do not exist. Accessible exemplars may influence group judgments in addition to, and not instead of, preexisting abstract knowledge. In a similar fashion, the fact that group judgments may be influenced by contextual variations does not indicate a lack of abstract knowledge.

3. It is important to note that this is not an argument about economy of storage. In this model, all encountered exemplars may be retained in memory. However, storing those exemplars and using them are different matters. The point here is that retrieving and summarizing large numbers of exemplars is less efficient and more easily disrupted than activating a stored abstraction.

4. There are certainly cases in which attribute-based impressions are as easily derived as stereotype-based impressions. There are also cases in which stereotype-based inferences depend more on bottom-up and less on top-down processes than attribute-based inferences. This has led some researchers to conclude that the distinction between stereotyping and individuation is not a useful one, and that stereotype use is not inherently more efficient than individuating (e.g., Kunda & Thagard, 1996). Although we concur that it is not necessarily the case that stereotyping will always be more efficient than individuation, we also believe that it just so happens to be true most of the time. This is the case for a number of reasons. For example, whereas stereotypes typically provide access to knowledge along a wide range of dimensions, many attributes do not generalize beyond their own specific meaning (Andersen, Klatzky, & Murray, 1990). Moreover, some kinds of individuating information (e.g., behavioral information) must be extracted in a bottom-up fashion, which is rarely the case for many stereotypes (particularly those associated with obvious physical characteristics). For a complete discussion of this matter, see Bodenhausen, Macrae, and Sherman (1999).