

5

Dual-Process Models

Nonviolent direct action seeks to create such a crisis and foster such a tension that a community which has constantly refused to negotiate is forced to confront the issue. It seeks so to dramatize the issue that it can no longer be ignored. My citing the creation of tension as part of the work of the nonviolent resister may sound rather shocking. But I must confess that I am not afraid of the word “tension.” I have earnestly opposed violent tension, but there is a type of constructive, nonviolent tension which is necessary for growth. Just as Socrates felt that it was necessary to create a tension in the mind so that individuals could rise from the bondage of myths and half-truths to the unfettered realm of creative analysis and objective appraisal, so must we see the need for nonviolent gadflies to create the kind of tension in society that will help men rise from the dark depths of prejudice and racism to the majestic heights of understanding and brotherhood.

—MARTIN LUTHER KING, JR., “Letter from Birmingham Jail” (April 16, 1963)

Woven into King’s (1963) eloquent jail-cell wake-up call to America are many of the key elements that contribute to modern-day psychological models of social cognition. As we have been reviewing in the early chapters of this book, people have a tendency to rely on “less effort” and categorical thinking when evaluating others. Thus the categories people rely on, and the lack of effort exerted in drawing conclusions about others, can lead them to apply faulty and inaccurate categories—to lean on stereotypes, schemas, and prior expectancies. This is what we might assume King meant when he described people as being held bondage by “myths and half-truths.” People often fail to arrive at the “unfettered realm of creative analysis and objective appraisal,” and in Chapters 3 and 4, we have examined why this might be the case. Whether it be due to limited capacity or to the fact that people prefer to “think less” and rely on the least-effort principle, they often fail to reach the heights of objective and creative analysis of others.

But King's assumption is not that people are incapable of breaking free of these bonds. Instead, people are seen as having the cognitive flexibility to rise above this simplistic way of thinking about others and to embrace an objective appraisal. This belief that people may typically rely on categorical thinking, but can exert the flexibility of will and mind to be more objective and effortful, is the basic premise of *dual-process models*. Thus, on the one hand, processes of social cognition involve a reliance on stereotypes, heuristics, categories, schemas, and "half-truths." Because in these instances people rely on general, preexisting theories about others, and from these theories make inferences about specific individuals, this type of processing is generally referred to as *top-down* or *theory-driven processing*. On the other hand, people can exhibit quite a good deal of "creative analysis" and "objective appraisal" in thinking about others. We have already seen evidence of this in Chapter 4, in discussing the type of mental work that is triggered when information inconsistent with an expectancy is encountered. It is not merely ignored and dismissed; in fact, significant cognitive work may be exerted toward resolving the inconsistency. In these cases, people do not rely on prior theories, stereotypes, and heuristics, but examine the qualities of each individual when constructing an impression. Dual-process models use several interchangeable terms to refer to this type of effortful and elaborate processing of information about an individual: *personalization*, *elaboration*, *systematic processing*, and *attribute-oriented processing*. More generally, because impressions are being built from the observed data (the behavior of another), such processing is referred to as *bottom-up* or *data-driven processing*.

A second ingredient in King's (1963) analysis is also of central interest to dual-process models: Namely, what is it that motivates people to shift from the comfort and ease of top-down processing to exerting the effort required of systematic processing? In King's words, the answer is that a tension in the mind motivates people to rise from the "depths of prejudice" to the "majestic heights" of understanding. Tension (psychological, not physical) forces them to analyze others carefully and reanalyze their own prior views. Psychologists have had much to say about the motivational powers of a state of psychological tension. Tension, as reviewed in the Introduction to this book, is a psychological state people find aversive and are motivated to reduce. In this case the tension arises from top-down thinking's being deemed unacceptable. This tension can be reduced through engaging in more elaborate processing of information relevant to the person who had previously been judged in a top-down fashion.

A final ingredient in King's (1963) letter also maps quite closely onto dual-process models of impression formation: Namely, what triggers tension in the mind of the perceiver and initiates the shift from schematic and stereotypic processing to individualizing and attribute-oriented processing? In King's approach, tension is initiated by the behavior of the person being perceived. If the person being perceived by others can act in a manner that is opposed to the prevalent biases, stereotypes, and expectancies, then it forces the people doing the perceiving to confront the fact that their theory-driven approaches to understanding this person are not working. Regarding the application of this idea to King's involvement in the civil rights movement in the United States, King chose to advocate nonviolent protest—a challenge to the existing stereotype that African Americans were more violent and less intelligent than Whites. By consistently challenging the stereotype, King believed that the members of the majority group (White Americans) would be forced to reevaluate their stance and think more deeply about the issues and about African Americans. Dual-process models recognize this as a viable form of tension induction. Consistent behavior that is clear, diagnostic, and counter to prevailing stereotypes (such as when a supposedly violent group of Black men, according to a

stereotype, responds to brute force with nonviolent protest and peaceful pleas for change) can motivate a shift away from stereotypic processing and can trigger elaborate and detailed thinking about an issue or person. Although this is not the only manner tension can be produced, King (operating without the benefit of the research we will review here) struck upon one good method available to him.

Martin Luther King, Jr., famously dreamed of a world where people are judged not by the color of their skin (categorical/stereotypic processing that is characteristic of being “theory-driven”), but on the content of their character (personalized and systematic processing that is characteristic of being “data-driven”). The latter requires perceivers to rise up from the “depths” of categorical thinking to a more elaborate and systematic analysis of another person as an individual. To do so requires that perceivers have the flexibility to shift from being theory-driven to data-driven. If capable of doing so, people are perhaps best viewed not as cognitive misers, but as *flexible interpreters* (Uleman, Newman, & Moskowitz, 1996) or *motivated tacticians* (Fiske & Taylor, 1991): They have a tendency to rely on preexisting theories, but can be motivated to utilize the flexibility of their cognitive system and think about others in a more elaborate fashion. As we will see, being more effortful in evaluating others may not guarantee being objective and fair in thinking about others. After all, people can exert considerable effort to uphold an existing bias against someone and to cling to stereotypes. Rationalizations are often a quite effortful way to maintain whatever beliefs people wish to maintain (e.g., Kunda & Oleson, 1995). But although exerting mental effort does not guarantee thinking objectively and rationally, dual-process models at least view people as having the capability to approach the objective appraisal of others of which King dreamed.

A SAMPLER OF DUAL-PROCESS MODELS

With enough time and testing of defining cues, “best fit” perceiving can be accomplished for most but not all classes of environmental events with which the person has contact. There are some objects whose cues to identify are sufficiently equivocal so that no such resolution can be achieved, and these are mostly in the sphere of so-called interpersonal perception: perceiving the states of other people, their characteristics, intentions, etc. on the basis of external signs. And since this is the domain where misperception can have the most chronic if not the most acute consequences, it is doubtful whether a therapeutic regimen of “close looking” will aid the misperceiver much. . . . But the greatest difficulty rests in the fact that the cost of close looks is generally too high under the conditions of speed, risk, and limited capacity imposed upon organisms by the environment.

BRUNER (1957, pp. 141–142)

Dual-process models describe people as having a default strategy in which inferences are formed and impressions are made via heuristics, schemas, stereotypes, and expectancies. This strategy at the effortless end of the information-processing continuum has been described in Chapter 4 as having evolved out of a need to manage the complexity of the environment using a cognitive system with limited resources. The restrictions to one’s cognitive resources can arise from conditions in the environment such as these: working on many tasks at once, which places people under cognitive load (“limited capacity”); working on tasks under a deadline, where time pressure limits how long you can deliberate about the qualities of others (“speed”); and working on tasks that, in and of themselves, are complex and difficult. Bruner asserts that interpersonal perception is an example of such tasks: Person perception is a domain where the thing being per-

ceived (another individual) is highly equivocal. Even with all the time and capacity in the world for effort and analysis, people may never come to truly understand the character of another person by observing his/her behavior. Thus they often simply rely on the strategy of using the cognitive shortcuts that their prior theories and categories supply. However, Bruner paved the way for dual-process models by asserting that people are sometimes willing to pay this cost and engage in the effort of what he called a “closer look” at the information. The important psychological questions concern when, why, and how they are willing and able to pay such costs.

Brewer’s Dual-Process Model of Impression Formation

The term *dual-process model* was coined by Brewer (1988) in describing a comprehensive theory of the processes involved in impression formation (depicted in Figure 5.1). Although theories that are now categorized as dual-process models had existed in social psychology prior to that point (e.g., Chaiken, 1987; Kunda, 1987; Neuberg & Fiske, 1987; Petty & Cacioppo, 1986; Trope, 1986a), Brewer’s model built an important bridge between research focused on how people use schemas, categories, and heuristics in judging others in a fairly effortless and automatic way, and much earlier research focused on the rational and methodical processes that people employ when forming

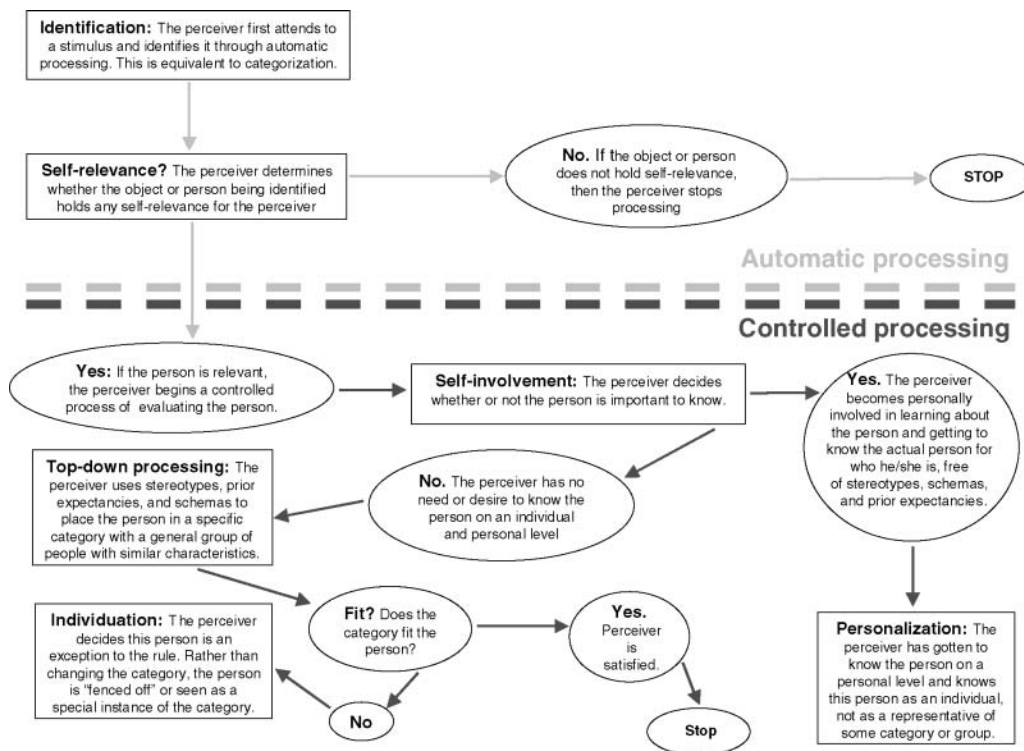


FIGURE 5.1. Brewer’s (1988) dual-process model of impression formation. From Brewer (1988). Copyright 1988 by Lawrence Erlbaum Associates. Adapted by permission.

attributions of others (reviewed in Chapter 6). For many years, it was known that people use each of these processes in their social cognition. However, the conditions that delineated when social cognition would be marked by an overreliance on categories and fairly automatic processes, and when it would involve an attention to details and fairly effortful processes, had not been established. Several integrative models emerged between 1986 and 1990 that attempted to specify the conditions governing when and why people shift from the more automatic type of processing to the more systematic type (e.g., Brewer, 1988; Chaiken, Liberman, & Eagly, 1989; Devine, 1989; Fazio, 1990; Fiske & Neuberg, 1990; Gilbert, 1989; Kruglanski, 1990; Kunda, 1987; L. L. Martin, 1986; Petty & Cacioppo, 1986; Trope, 1986a; Wyer & Srull, 1989). These models share a set of assumptions in which impression formation is organized according to stages that occur sequentially, such that the “miserly” perceiver will not expend resources for further processing unless certain conditions are met.

The attempt to integrate the two types of cognitive processes that had been the focus of impression formation research helped to reintroduce the notion of motivation—a concept that had been diluted (though not absent altogether) in the cognitive revolution—to mainstream social psychology. In research studying the automatic nature of these processes, goals were often held constant, rendering them almost obsolete. In research of the 1960s and early 1970s on attribution (see Chapter 6), motivation was usually treated as an explanation for why people are biased and why these typically rational processes break down. Rather than being seen as central to the process, it was viewed as sending the process awry. However, when researchers began attempting to integrate different forms of thinking within a given perceiver, they soon rediscovered that motives and goals play a central role requiring that they be integral in any model of social cognition. The first volume of Sorrentino and Higgins’s (1986) *Handbook of Motivation and Cognition* helped to awaken the field to this issue, and dual-process models soon sprung forth. As a case in point, Brewer’s (1988) dual-process model views the role of motivation as central to determining the type of social cognition perceivers will engage in: “The majority of the time, perception of social objects does not differ from nonsocial perception in either structure or process. When it does differ, it is determined by the *perceiver’s* purposes and processing goals, not by the characteristics of the target of perception” (p. 4; emphasis in original).

Identification

Brewer’s model specifies that both automatic and controlled processing contribute to our ultimate impression of another person. At the automatic level, the model begins with processes of attention and identification not unlike the silent process of categorization that Bruner (1957) described and that we have reviewed in Chapter 3. When a person enters our social environment, we perceivers must first focus attention on that person and identify the person as being present, as having certain features, and as having performed certain types of behavior. Brewer (1988) referred to these automatic processes that allow us to capture the presence of a target stimulus (in this case, a person) and identify the basic features that the target is eliciting as *identification*. Thus, without any conscious effort by us as perceivers, some form of information is perceptually salient enough to cue a social category. The model suggests that a few social categories are used often and consistently enough to be triggered automatically (e.g., age, ethnicity, race, and sex). Exactly which of these categories becomes the superordinate cate-

gory label that may organize subsequent impression formation and processing depends on features of the context, the perceiver, and the target. Fiske and Neuberg (1990) assert that qualities that “possess temporal primacy, have physical manifestations, are contextually novel, are chronically or acutely accessible in memory, or are related in particular ways to the perceiver’s mood will tend to serve the role of category label” (p. 10).

Sometimes the categorizations we make by using this process of matching observable features of a target to our existing mental representations of “person types” are wrong. We may assume that a person who speaks fluent English without accent is an American, when in fact the person is Cuban; we may assume that a person with fair skin is not an African American, when in fact many people with one White parent and one Black parent identify themselves as African Americans regardless of their skin color. As Allport (1954, p. 20) says, “Sometimes we are mistaken: the event does not fit the category. . . . Yet our behavior was rational. It was based on high probability. Though we used the wrong category, we did the best we could.” For reasons elaborated in Chapter 3, categorizations are an essential and functional part of our understanding of the world, and this feature-matching process allows us to label and type each stimulus we encounter in a fast, efficient, economical, and relatively (if not 100%) accurate way.

Determining Relevance

The primary outcome of this initial identification stage is a preconscious decision as to whether further processing is necessary. Our categorization will allow us to detect whether the person is relevant or irrelevant to us. If a person we have encountered is irrelevant to our current goals and purpose (e.g., a stranger passing us on the subway platform, a driver stopped at a red light as we walk down the street, a clerk at the market placing Ring Dings on the shelf), then further processing, and the expenditure of cognitive effort to think about this person, are not necessary. Impression formation processes are halted. However, if the person is deemed somewhat relevant—and this can range from a minor degree of relevance (such as a store clerk who is placing Ring Dings on the shelf when we want Ring Dings) to a major degree of interest (such as a stranger passing us on the subway platform who attempts to mug us)—then impression formation processes are triggered. At this point, a critical choice is made between the two alternative processing modes described in the model. Thus even the initial, preconscious decision that determines whether we dedicate any energy to thinking about a person and process further information about the person depends on goals (i.e., the relevance of the target).

Once a target has been deemed relevant, processing that is automatic (preconscious automaticity) is no longer engaged. There is a shift toward processing that fails to meet all of the criteria for an automatic process, if only because such processing is goal-dependent. As discussed in Chapter 2, any processing that requires a goal for its initiation, as opposed to the mere presence of a stimulus in the environment, is not considered automatic. This is not to say that such processing cannot occur with extremely little effort, or even without our awareness. But processing from this point onward does require a goal—the goal of needing to know more about the person because of his/her potential relevance to us. The degree to which conscious effort and awareness are engaged by subsequent processing is now what distinguishes the subsequent stages of information processing and establishes the two broad modes of processing (categorization vs. personalization) that comprise the dual processes in the model.

Categorization/Typing

Motives do not only determine whether we proceed to think about others; they also determine *how* we proceed to think about them. The decision to gather further information about a person does not mean that it will always be an effortful, deliberative, rational, and objective process. We can produce judgments of others in a categorical way, relying on schemas, heuristics, and prior expectancies about the category of people the particular individual being judged belongs to (e.g., stereotypes). Brewer (1988) has referred to this relatively effortless type of thinking, which is anchored by and directed exclusively by the categories that have been (rather passively) triggered, as *categorization* or *typing*. In this type of impression formation, our judgments about a person are based on available “person types” or schemas that are matched to the information at hand about the person being evaluated. An iterative, pattern-matching process is conducted until an adequate fit is found between one of the person types or categories stored in memory and the stimulus characteristics. The process starts with a rather general category, and if there is no match, subtypes of the general category that provide more specific types are compared against the target person. Brewer offers a concrete example of an abstract superordinate category and the more concrete subtypes within that category. The superordinate category “older men” is fairly abstract and contains a wide range of features. Within that superordinate category is the more specific category type “businessmen,” which is still relatively abstract, but less so than the category label in which it is nested (“older men”). Moving downward in the category hierarchy, we arrive at even more specific subtypes of the category “businessmen” that have fine-grained sets of features, such as an “uptight authoritarian boss who is a tightwad and a stickler for detail” (Brewer, 1988, p. 12).

Given these different levels of category types, Brewer has concluded that the exact manner in which a given person is categorized depends not only on the features of that person being matched against a category, but on the level of abstraction within the category structure the feature matching process begins. Exactly where in a perceiver’s category system this process begins can be determined by several factors: what categories are automatically triggered during identification, cues in the current situation, and the processing goals of the perceiver. For example, if a person is initially classified according to age as a teenager, and then information is attended to that indicates a counterstereotypic role (e.g., the teenager is a writer for *Rolling Stone* magazine), then his/her role as a writer will probably be organized as a specific type of teenager (e.g., a “whiz kid”). However, if age is not salient in the initial categorization—for example, if the editor of *Rolling Stone* magazine called the manager of the band Led Zeppelin and said he could expect Cameron (an age-unspecified writer) to come and tour with the band—then the initial impression will start with a more general category, “writer,” rather than “whiz kid.”

These categorizations are important for several reasons. First, category-driven impressions might be the sole contributor to one’s final judgment of the person, constituting the entire basis for one’s opinion. That is, there are times when perceivers are perfectly happy to rely on their categories to supply them with broad over-generalizations that can extend from the category level to individual members of the category. Allport (1954) provides an excellent example of this when discussing how stereotypes of groups can be extended to an individual member of a group once he/she is categorized by the perceiver as a member of that group. The very first sentences of

Allport's seminal book make this point: "In Rhodesia, a white truck driver passed a group of idle natives and muttered, 'They're lazy brutes.' A few hours later he saw natives heaving two-hundred pound sacks of grain onto a truck, singing in rhythm to their work. 'Savages,' he grumbled" (p. 3). Second, even if the impression one forms evolves and is transformed in later stages of information processing, the category-based judgment will still serve as an anchor and input (starting point) for any subsequent judgment (e.g., Park, 1989). Finally, categorical processing not only serves as an anchor around which future, effortful judgments shift; it also anchors the effortless processing that occurs during categorization. That is, during the process of trying to match "person types" against a given person's characteristics in order to categorize him/her, the initial categorization, even if not successful at categorizing the person, has an impact on how the subsequent feature mapping and person typing proceed. If the initial match between a category and a person's features is not successful, this first attempt will not be discarded for a new, perfectly irrelevant schema/person type to match against the target person. Instead, if the initial fit is not adequate, "the search for an appropriate categorization will be directed downward, among subtypes of the original category, rather than horizontally among alternative categories at the same level of abstraction. . . . Thus, initial category activation sets in motion an iterative process that constrains the final category selection" (Brewer, 1988, pp. 18-19).

Individuation and Personalization

The processes of categorization seem to involve no more motivational impetus than trivial levels of self-involvement. If a person is somehow relevant to us, category-based processing is initiated. But as our needs and goals as perceivers dictate, we can engage in more detailed, elaborate, and effortful types of cognitive processing. This is not to say that all category-based processing must be relatively effortless. At times we expend quite a good deal of energy and effort to maintain our categories. Chapter 4 has discussed how we process information that is inconsistent with (or incongruent with) an expectancy or category that has been triggered; we often work quite hard to continue to process in a categorical way (because, despite this effort, it is still less effortful than reorganizing our whole category structure and changing an entire mental representation). Brewer (1988) provides another example of fairly effortful, category-based processing, in a process labeled *individuation*. In this type of processing, information we receive about a person that is inconsistent with the category label is not disregarded; nor is a more specific subtype of the category called into play to describe the person. Instead, the individual is treated as an isolated case or a specific instance of the particular category. Thus the qualities of the individual are processed in detail—not for the purpose of making that individual fit well with the existing features of the category, but to create a specific and detailed novel instance of the category. Brewer provides a good example: "The first anchorwoman that appeared on a national television news program was no doubt highly individuated as a member of the category of news broadcasters" (p. 21).

There are times, however, when no form of category-based processing is deemed appropriate by us as perceivers. It is at these moments that the model suggests a wholly different form of processing in both format and organizational structure. At this point in the process, the individual becomes the basis for organizing information. Here the category label is no longer superordinate, and the individual's attributes serve as the organizing framework for understanding that individual. Brewer referred to this type of

information processing as *personalization*. Brewer (1988, p. 22) asserts that “the new organizational structure is a sort of mental flip-flop of the category-based structure,” and provides an example via analysis of the statement “Janet is a nurse.” This information can be organized and analyzed in a category-based way, with Janet being treated as a type of the category “nurse.” Alternatively, it can be organized and analyzed in a way that treats “nurse” as subordinate to, and a feature of, the properties of an individual named Janet. In one case, the category “nurse” is the organizing structure around which the emerging mental representation is built; in the other case, it is the individual, Janet. “The concept of ‘nurse’ as a feature of Janet would contain only those aspects of nursing that are characteristic of Janet in that role, and would be disassociated from the prototypic representation of nurse as a general category, which may contain many features not applicable to Janet” (p. 22).

Exactly what sort of conditions lead us as perceivers to conclude that no form of category-based processing is appropriate? Brewer (1988, p. 25; emphasis in original) asserts that “personalization implies some degree of *affective* investment on the part of the perceiver, either because of the target person’s relationship to the perceiver’s personal or social identity, or because of the target’s relevance to the perceiver’s personal goals.” Thus, if someone is similar to us, we are less likely to rely simply on category-based processing and more likely to expend the mental energy to personalize him/her. Similarly, if a person is linked to important goals that we seek to attain (e.g., a coworker whose work on a joint project will also determine how we are evaluated, a manager whose evaluation of us will determine whether the company decides to promote us), we are more likely to move beyond simple category-based processing and form impressions of this individual based exclusively on the attributes and features we have taken the time to detect.

The Heuristic–Systematic Model

The heuristic–systematic model (HSM; Chaiken et al., 1989) provides a fairly comprehensive account of the mechanisms that drive social cognition. According to the HSM, people’s thinking about themselves and others can be described in terms of two broad information-processing strategies. These two strategies are often described as endpoints on a continuum of processing effort that can be exerted in forming judgments (see Figure 5.2).

Two Processing Modes

At one end of the continuum is the relatively effortless *heuristic processing*—a reliance on prior knowledge such as schemas, stereotypes, and expectancies that can be imposed on information so it will easily fit into existing structures. Instead of exerting effort toward thinking carefully about content (whether it be the content of a person’s character or the content of a message), people instead preserve their processing resources by relying on superficial assessment. That is, people rely on heuristics, shortcuts, and rules of thumb in thinking about their social world. Such rules of thumb, categories, and expectancies are learned through experiences in the social world, and are readily called to action when people encounter a cue in the environment that signals the relevance of one of these previously learned rules or responses. As an example, past experience indicates that a person whose point of view achieves a wide consensus among others is often a person whose point of view is correct. Armed with this “Consensus implies correct-

There are two general strategies for processing information about people. These can be conceived of as opposing ends of a continuum. The continuum represents the amount of mental effort one exerts when thinking about others. Many believe the human default is to exert less, rather than more, effort.

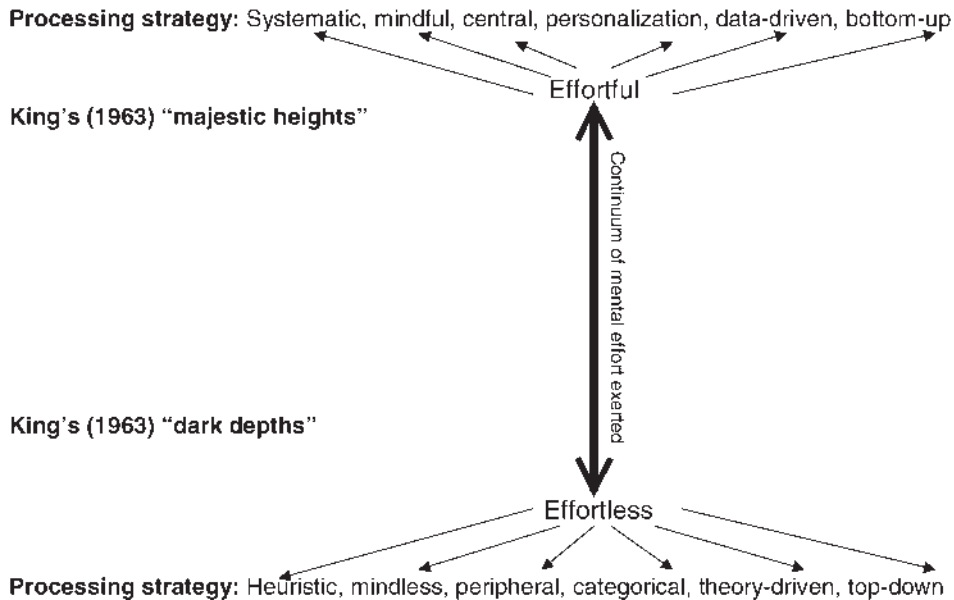


FIGURE 5.2. Two processing modes.

ness" heuristic, perceivers may react to future encounters with a majority position on an issue in a heuristic fashion. They can superficially assess the situation by relying on the rule of thumb, forming an opinion in a top-down fashion without scrutinizing the actual arguments used by the majority in advocating their position. Thus, if most people agree with something, these perceivers will agree with it too, using the consensus as the basis for their opinion.

On the other end of the continuum is *systematic processing*. Systematic processing involves a detailed evaluation of the qualities and behaviors of others, as well as the reexamination of personal thoughts and prior beliefs about the stimulus. Individuals will attempt to formulate a judgment through reconciling the current data that they have analyzed quite closely with their reexamined beliefs, using processes such as integration (Kelman, 1961) and validation (Moscovici, 1985b). Because perceivers are actively thinking about how some stimulus information relates to other knowledge they possess, systematic information processing requires a sufficient amount of their limited cognitive capacity or processing resources in order to be carried out, as well as some degree of motivation to instigate such effortful processing.

Consider the example used above. If past experiences indicate that a person whose point of view achieves a wide consensus among others is often a person whose point of view is the correct one, then perceivers develop expectancies and heuristics not only about what it means to be in the majority, but also about what it means to be in the minority (Moscovici, 1985a). Armed with these "Consensus implies correctness" and

“Lack of consensus implies incorrectness” heuristics, people may react to future encounters with both a majority and a minority position on an issue in a heuristic fashion. In particular, they may simply react with disdain and disregard for what the minority has to say. However, as was the case with the nonviolent protests of the civil rights movement in the United States, the perceivers can be shaken from “top-down” thinking so that they no longer are imposing their previously held theories, heuristics, and biases. If people doubt the reliability and validity of their heuristic thinking (the stereotypes about the minority seem out of place or incorrect), this will trigger thought processes whereby the minority’s position is deeply evaluated and thoroughly appraised, and their own positions (and biases) are reevaluated. The perceivers will examine the data, building a judgment from the bottom (the facts), rather than imposing one from on high.

This assumption of two processing modes raises two further questions: Why assume that one of these is the default mode, and what are the appropriate motivational conditions that lead one to shift from the default to the other mode of processing? The HSM presents two principles that answer these questions—the least-effort principle and the sufficiency principle. Each should be familiar by now, as they reflect principles already discussed earlier in the book.

The Least-Effort Principle

Whereas Allport (1954) put forward the principle of least effort (reviewed in Chapter 4), the HSM puts forward the least-effort principle! As you might expect, these are one and the same idea. Recall that the least-effort principle asserts that when performing cognitive tasks, people prefer less mental effort to more mental effort. This suggests that the default processing strategy will be the one requiring the least amount of effort and usurping the least amount of capacity—the heuristic route. How can we detect whether the use of heuristics is the default way of processing information? By examining whether heuristic processing is found to occur even when there are no specific motivating circumstances to engage in this processing mode. Such processing should merely require the presence of the appropriate cues in the environment to trigger the use of the heuristic. For example, if perceivers are evaluating the positions of a minority by relying on heuristics, the recognition that the minority lacks any consensus should be a cue that will trigger relevant rules of thumb, which the perceivers then will rely on if not motivated to think “harder.”

Petty, Cacioppo, and Schuman (1983) provided an illustration of people using heuristics as a default strategy. Their research examined why people are persuaded by celebrities who endorse a product in a commercial. The logic was that seeing a celebrity leads people to rely on simple rules of impression formation, such as “Famous people are trustworthy” or “Attractive people know a lot about dietary and cosmetic products,” rather than evaluating the details about the product. Petty and colleagues had research participants read about a razor that was being endorsed by either a sports star or an unknown person. There were two potential sources of influence in this situation. Participants could read the text of the endorsement carefully and evaluate the product based on its merits, or they could superficially evaluate the product and rely on the word of the endorser. If participants were using the former strategy, they should find the product equally good, regardless of whether the celebrity or the “regular guy” was the person endorsing it. If they were relying on heuristics, the celebrity should be much more persuasive at selling the product than the “regular guy,” even if the text of their sales

pitch was the same. The results revealed that rather than using the text to guide impressions, the participants did indeed use a heuristic about the person delivering the text. Thus instead of evaluating people and products in detail, scrutinizing their qualities, people rely on simple and easy rules.

The Sufficiency Principle

The sufficiency principle asserts that for whatever task people are confronted with—whether it be forming an impression of someone, planning how to act toward someone, forming an attitude, making a decision, or simply comprehending some information—there is a point at which they feel that their task is completed and they can move on to the next task at hand. This point is said to be achieved when the individuals feel *confident that they have sufficiently performed the task* that was set before them. This point of sufficient confidence that allows for a feeling of task completion can be conceived of as a threshold—a *sufficiency threshold*. That is, in almost all things, people desire to feel somewhat confident about their thoughts, feelings, and actions. And when they are gauging their sense of confidence in the validity of their own beliefs, attitudes, and actions, this can range from not at all confident to extremely confident. People may not (in fact, it is likely that they typically do not) always desire to have extreme confidence in what they say and do. They are sometimes happy simply to have moderate amounts of confidence in their judgments and feelings, so long as they feel it is *enough* confidence to warrant relying on those judgments/feelings in their interactions.

This sufficiency principle allows us to conceive of the point at which we perceivers feel confident enough in our thoughts and actions as a threshold that can be set anywhere we choose, and as a point that is subject to change from situation to situation. For example, when we are evaluating a TV commercial or a stranger at the market, we may feel relatively confident in the opinion we form, despite barely listening to what is being said to us. Yet when evaluating a love interest, we may cling to the person's every word with an effort and detail that we might not have known we possessed. We seek to feel extremely confident that our responses to this person are accurate and appropriate, so we set the threshold for deciding when our thoughts and deeds are sufficient quite high.

The Relationship between Effort and Sufficiency

The sufficiency principle asserts that although people prefer least effort, they must exert *enough* effort to reach the sufficiency threshold. If people desire the feeling that their judgments are good enough, they will only rely on heuristics in forming judgments if those heuristics deliver to them a sense of *judgmental confidence*. For the most part, heuristics provide sufficient responses to allow people to navigate through social interactions. The *product* of heuristic processing is a judgment based on relatively little effort. To the degree that relying on little effort yields a sufficiently good product—judgments and behaviors that leave people feeling sufficiently confident that they have acted in a reasonable manner—then they are happy to rely on such heuristics and do not need to engage in any further processing of information. If heuristic processing is experienced as producing inadequate actions, judgments, or feelings—that is, the level of confidence in the judgments that heuristics produce falls short of the threshold—people will exert more effort and continue working on the task at hand until a feeling of sufficiency is achieved and the threshold is surpassed. Thus the least-effort principle

exists because using less effort (relying on heuristics) typically gives people a sense that they have formed judgments in which they are sufficiently confident (see Figure 5.3).

For example, experience may teach one that people who are politically conservative are against abortion. If this position opposes one's own feelings, one may try to avoid discussing the topic with a conservative who attempts to discuss it. A cue—the other person's conservatism—triggers a heuristic in the mind (“Conservatives have extreme views, so don't bother arguing with them; they will not yield”). The heuristic, in this case, works. It allows one to have inferences that seem reasonable enough, and allows one to plan behavior that will yield a smooth and trouble-free encounter. Thinking less has produced sufficient inferences to generate successful social behavior. However, if one's confidence in relying on this heuristic is not sufficient, one will be forced to increase processing effort; that is, when one encounters a conservative, one may deem a sole reliance on heuristics to be inappropriate. It may be true that conservatives typically oppose abortion, but this given individual may have been expressing a dissenting, pro-abortion view that one's heuristic processing did not allow one to detect. One must initiate more elaborate forms of evaluating this person until one has produced a judgment that one feels is sufficient (is reasonable enough to rely on). At what point is effort frozen (rather than continually increased)? When one's current confidence in a judgment has catapulted over the sufficiency threshold.

In this hydraulic-like model, the energy that drives the shift from effortless to effortful thinking about other people is a pressure to have greater accuracy and to remove feelings of doubt. If least-effort processing produces judgments, behaviors, and evaluations that perceivers have relatively little confidence in, a problem exists. They desire a certain level of confidence in the cognitions that their processing system produces (this is said to be the sufficiency threshold). When relying on “least effort” yields

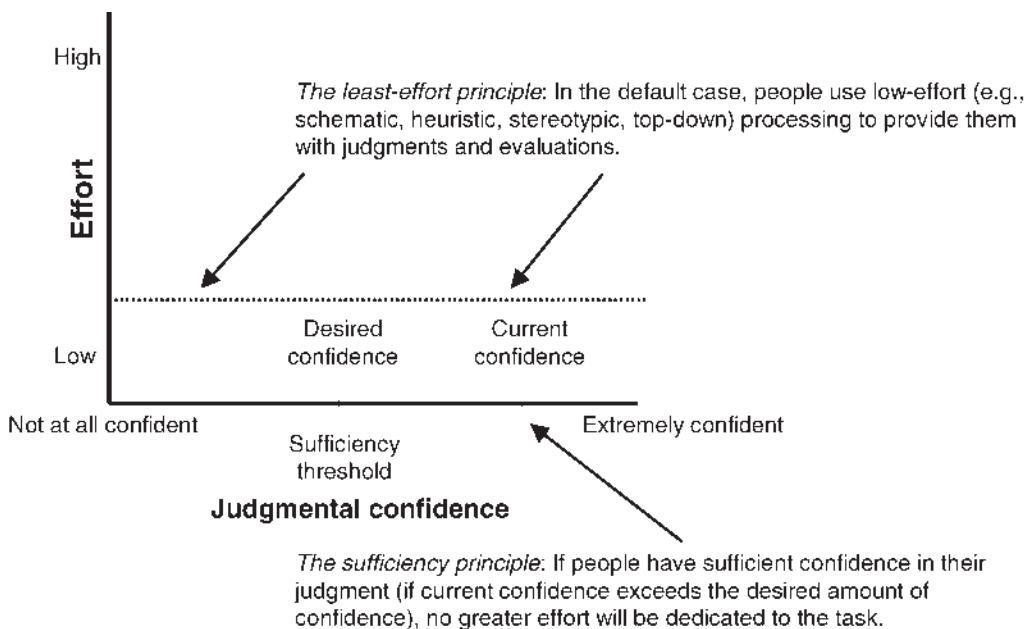
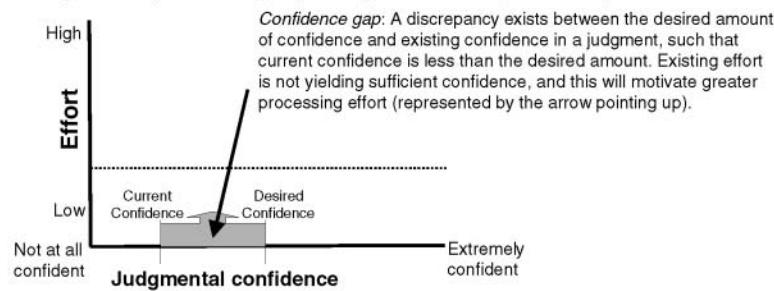


FIGURE 5.3. The least-effort principle and the sufficiency principle.

judgmental products that are not sufficient—that fall below the sufficiency threshold—a *confidence gap* is said to exist (see Figure 5.4A, Time 1). This is said to be experienced by the perceivers as an uncomfortable and displeasing state of psychological tension. Because the tension state is not pleasant, the perceivers are motivated to reduce the state and eliminate the gap between their current level and desired level of confidence. How can this be accomplished? The pressure from the experienced confidence gap pushes processing effort upward. Because of the hydraulic-like relationship between confidence and effort, lacking confidence in judgments generates increased effort in the processes that yield or produce these judgments. Thus the confidence gap can be reduced as one makes the transition from the least-effort principle's reliance on minimal processing effort to expending more effort. Perceivers will now process information in a more systematic fashion (see Figure 5.4).

In Figure 5.4B, the left-hand side of the figure illustrates how an initial increase in effort (as represented by the gray arrow) can lead to a reduced confidence gap (as represented by the black arrow). However, in this instance the moderate amount of increased effort is not sufficient to produce a judgment in which the person feels “confident enough.” Existing confidence at this point (Time 2), though higher than at Time 1, is still below the desired amount of confidence (the sufficiency threshold). Therefore, the individual still experiences a confidence gap and the pressure to reduce this gap. This, once again, is accomplished through exerting even greater processing effort

A. The cognitive system experiencing psychological tension (at Time 1).



B. Processing effort is increased to reduce psychological tension (Times 2 and 3).

Systematic processing: Effort is expended to elaborate on a stimulus until judgment is “good enough.” At left, the effort exerted at Time 2, though increased from Time 1, produces a confidence level at Time 2 that is not sufficient; a confidence gap (though reduced from that existing at Time 1) persists. At right (Time 3), effort increases and confidence in judgment is now sufficient.




FIGURE 5.4. Dual-process models and the relationship between effort and confidence.

until the products of this data-driven analysis of the stimuli are finally recognized as sufficiently good. The increased effort has yielded an increase in confidence, such that at this point (Time 3) the processing effort has yielded judgmental products that surpass the desired amount of confidence (as evidenced on the right-hand side of Figure 5.4B by the dashed black arrow moving past the sufficiency threshold).

The Motivation to Make the Transition from Less to More Effort

The HSM suggests that the impact of heuristic processing on judgments will be greatest when motivation for systematic processing is low. Some trade-off must be struck between the goal of exerting least effort and the goal of having judgments one is sufficiently confident in. Although heuristic processing more fully satisfies the least-effort principle, systematic processing is more effective at producing greater amounts of confidence and is thus better able to satisfy the sufficiency principle. The assertion that humans process heuristically by default suggests that on balance, the trade-off between less effort and sufficiency is adequately met by heuristic processing; the sufficiency threshold is set low enough that it can adequately be reached by heuristic processing alone. Determining when people shift from heuristic to systematic processing revolves around the issue of when heuristic processing is *insufficient*. When the motivation for systematic processing exists, the impact of heuristic processing on judgments will be somewhat attenuated. In these instances, judgments based on information gathered through a deliberate and effortful strategy should be more reliable and trustworthy than those produced by a reliance on simple decision rules.

The next logical question arising from this discussion is this: What motivational circumstances are required to instigate systematic processing? We have already provided one type of answer to this question. The shift to systematic processing will be dependent on the existence of a gap between actual and desired confidence. This confidence gap is then experienced as an unpleasant tension state: People are now motivated to reduce the tension through systematic processing. However, the best answer to this question requires identifying *how* a confidence gap is produced or *how* to get people to experience their current judgments as insufficient. To alter what is experienced as “sufficient,” to experience a confidence gap, and to have motivation to engage in systematic processing, one of two things must happen. First, people’s confidence can be undermined, so that a judgment that was once firmly accepted is called into doubt. Here, the sufficiency threshold itself does not change; people’s existing level of confidence in their previously accepted heuristic changes so that it slips below the sufficiency threshold (e.g., Maheswaran & Chaiken, 1991). Second, desired confidence can be raised by a shift in the sufficiency threshold—a raising of the amount of effort needed in order to feel confident.

Undermining confidence. Confidence in a judgment (or attitude or belief or action) is not a fixed entity; it can change. At one point in time people might be confident in their use of a stereotype or a heuristic, only to feel unsatisfied with the judgments based on such minimal processing at a later point in time. In fact, this can occur (1) without a desire for any greater confidence than before, and (2) when the amount of confidence desired is quite low. For example, imagine you are meeting a woman, Jane, on Monday morning, and she shows up 40 minutes late without an explanation. You might feel confident in your opinion that Jane is unreliable, even if that opinion is based on a heuristic (such as “People who are extremely late for an appointment are unreliable”). In fact,

you might have a very low sufficiency threshold that allows you to be perfectly satisfied with this conclusion, despite giving minimal thought to evaluating Jane. However, on Tuesday afternoon you find out from a mutual friend that Jane is poor, and is embarrassed that she cannot afford a car or a watch. And these issues related to her being poor are what made her late (and to fail to tell you why she was late). In this case, you may no longer be confident in your heuristic-based judgment that Jane is unreliable.

Note that in this example, it is not the case that you desired a judgment that was based on better processing, or that you desired greater confidence than your original heuristic provided. You were initially perfectly satisfied with using little processing effort and relying on your heuristic. It was that the confidence you had in the conclusion your heuristic provided for you was later undermined, and your confidence then fell below the sufficiency threshold; your conclusion no longer provided a sufficiently good explanation for Jane's behavior (see Figure 5.5). Jane might be unreliable, but you are no longer confident enough in this heuristic-based conclusion, because new information (Jane's poverty caused her lateness; her pride caused her failure to explain this) has introduced doubt and lowered your confidence in the conclusion. As Maheswaran and Chaiken (1991, p. 15) have stated, it is "possible to promote systematic processing even when sufficiency thresholds are low by undermining perceivers' confidence in their heuristic-based judgments."

To illustrate this mechanism through which people are motivated to process systematically, an experiment will next be reviewed—one illustrating that systematic processing is initiated when conclusions drawn from a heuristic are subsequently challenged. Thus, while heuristic-based judgments and conclusions may be initially experienced as valid and sufficient (at Time 1 in Figure 5.5), a subsequent event can shatter confidence in these conclusions, lowering such confidence so that it falls below a personally defined threshold of what constitutes a sufficiently good judgment (at Time 2 in Figure 5.5). This triggers a switch to systematic and effortful processing that was previously not necessary (at Time 3 in Figure 5.5).

Maheswaran and Chaiken (1991) examined this point by looking at how consumers respond to advertising messages. Participants in their experiment were asked to read about and evaluate a new product, the "XT-100" answering machine. The participants first received information about how the product had performed in prior consumer tests; they were told that a sample of consumers who had been previously surveyed liked the product very much ("81% of 300 western consumers who had used the XT-100 were extremely satisfied," p. 16). Thus these participants in this experiment were now probably ready to evaluate the product fairly positively, regardless of what information they were actually given about the product. However, there was one proposed catch to this logic. If it turned out that the information describing the product was fairly negative, this would undermine participants' confidence in the heuristic. Most people were said to like the product, yet the information participants received about it would be highly incongruent with this claim. This unexpected information, inconsistent with the heuristic, should induce doubt and uncertainty in participants who were simply relying on the heuristic. Confidence in the heuristic should be lowered so that it fell below the sufficiency threshold. Contrast this with what should happen when participants received fairly positive information about the XT-100. In this case, they had an expectancy that the product will be good; they were predisposed to like the product because their heuristics informed them that it would be a good product; and a superficial assessment of the information provided (minimal processing effort) should inform them that all of this was correct—the product was fairly positively reviewed. Thus, in this case peo-

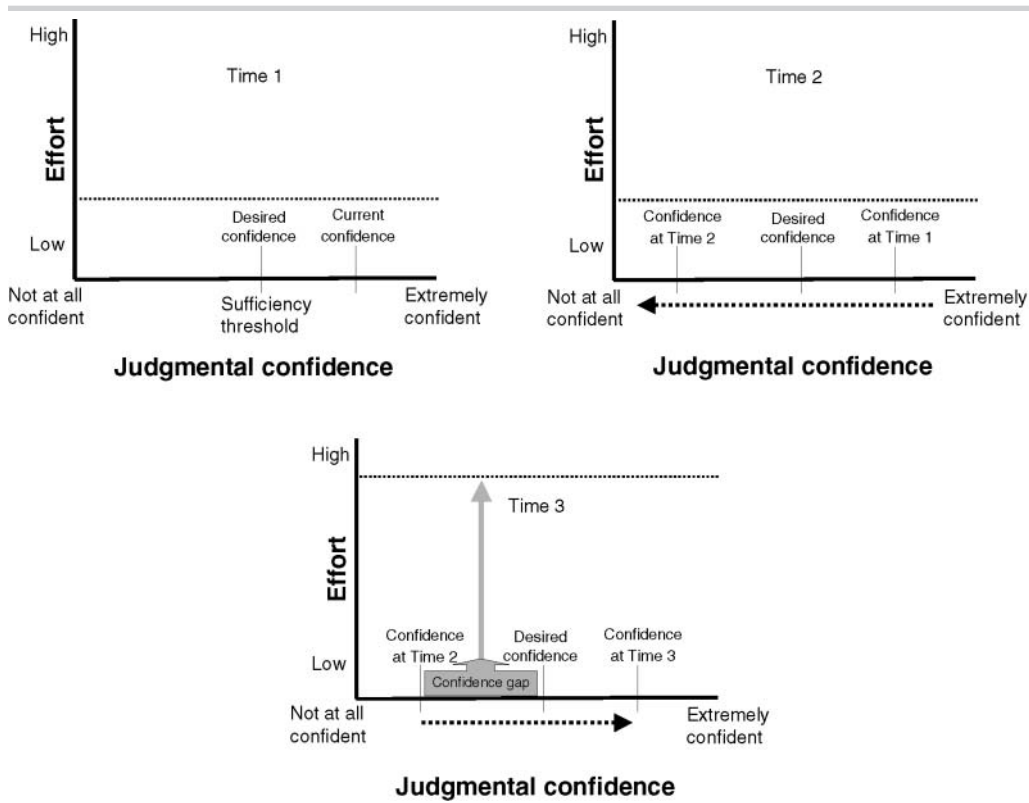


FIGURE 5.5. Instigating systematic processing by undermining judgmental confidence.

ple should have no need to exert too much effort in evaluating the product, relative to the people who were led to lack trust or confidence in their heuristic-based processing.

Maheswaran and Chaiken (1991) examined these hypotheses by manipulating whether participants received information about the XT-100 that was consistent or inconsistent with their heuristic-based expectancies. The information was said to have been provided by a product-testing agency's report that compared the XT-100 with competing brands. People receiving congruent information (a positive report on the XT-100) should feel free to process heuristically and exert relatively little processing effort when doing their evaluation of the product. People receiving incongruent information (a negative report on the XT-100) should lack confidence in heuristic processing and engage in more systematic processing (exerting greater processing effort). How could the researchers detect whether people were using heuristic versus systematic processing, exerting little versus much processing effort? The report had detailed information about how the product performed, as well as quite general, overall summary statements about the quality of the product. The logic of this experiment hinged on the idea that if people were systematically processing the message, they would pay much more attention to the details in the report, thinking about the XT-100's performance on these specific dimensions. Heuristic processors would focus instead on more general thoughts that summarized the overall quality of the product.

The results provided support for the notion that when confidence in a judgment

falls below a sufficiency threshold, systematic processing is triggered. First, people who received the unexpectedly negative report about the XT-100 reported having significantly less confidence in their rating of the XT-100 than people who saw a report that coincided well with the consensus information. Second, when people were asked how much confidence they *desired* to have, people who received the unexpectedly negative report about the XT-100 and people who saw a report that coincided well with the consensus information both expressed similar amounts of desired confidence; their sufficiency thresholds were equal. Combining these two points, the groups had equal amounts of confidence they desired to have in their judgment, and the members of one group (the people with confirmed expectancies) had an actual level of confidence that was sufficient in that it did not differ from the confidence they desired. The members of the other group, however, desired a level of confidence they lacked (because of the inconsistency between the report and the consensus of the people surveyed). Finally, the most important question of the experiment addressed what type of information processing was used. As predicted, people who lacked sufficient confidence in their judgment engaged in more thinking about the qualities and attributes of the product being evaluated. They gave more attention to and thought more explicitly about the details of the report than people who had sufficient confidence. They were also better able to remember these details when unexpectedly asked to try to recall the details of the report at the end of the experiment. Lowered confidence in the product of their heuristic-based processing triggered a confidence gap that motivated people to process systematically, despite the fact that the amount of confidence they desired was fairly stable. Actual confidence merely slipped below this sufficiency threshold and motivated people to think more deeply.

We have just seen how information incongruent with people's expectancies can lead them to alter the manner in which they think about and evaluate a novel and relatively uninteresting product (no offense to the XT-100). To what extent can these same ideas be moved outside the context of measuring attitudes toward objects and into the domain of person perception? When perceivers think heuristically and stereotypically about others, can their confidence in those stereotypes and heuristics be undermined by exposure to information contrary to the stereotype/heuristic? If so, is this sense of insufficient confidence enough to alter the manner in which perceivers think about the people they interact with?

Chapter 12 reviews one line of work examining what is called *minority influence*, which is focused on the perceptions of, attributions toward, and attitudes formed toward members of minority groups (and the strategies used by such groups to alter people's heuristically based perception, attributions, and attitudes). Over the past 30 years, there has been a growing body of empirical support for the idea that a minority that expresses opinions and behaviors incongruent with how the majority group expects others to think and act is able to force the majority to reexamine the expectancies previously held (e.g., Baker & Petty, 1994; Bohner, Erb, Reinhard, & Frank, 1996; de Dreu & de Vries, 1996; Moscovici, 1980, 1985a; Moskowitz, 1996; Mugny, 1975; Nemeth, Maysel, Sherman, & Brown, 1990). Behavior incongruent with what perceivers expect, especially from members of a minority group, initiates an alteration in the thought processes of the perceivers. Similarly, the "person memory" literature reviewed in Chapter 4 has illustrated how exposure to information about others that is incongruent with prior expectancies leads people to think more deeply about those others and engage in more elaborate processing effort. The dual-process notions we are reviewing here allow us to examine the question of why and how incongruent behavior

observed in a minority causes perceivers to shift from effortless to effortful thinking—it creates a confidence gap in which judgmental confidence is undermined. Chapter 12 provides a more detailed review of minority influence as being caused by the undermining of judgmental confidence. The basic point is that behavior from a minority that is inconsistent with the majority's expectancies relating to that minority will trigger systematic processing. This increased level of analysis creates the possibility that perceivers' confidence in relying on heuristic processing will be undermined, thus opening the door for the perceivers to change the way they think about that minority.

In sum, one way to trigger systematic processing, and to shake perceivers from the comfort of thinking heuristically and effortlessly about others, is to present those perceivers with information (behavior) that is so inconsistent with their prior expectancies that they can no longer rely on those expectancies alone. Such behavior removes the perceivers' ability to ignore or superficially assess the stimulus, and undermines the confidence they once held in the judgments yielded by their stereotypes and heuristics. This basic tenet of dual-process models was offered decades ago in Kelley's (1973) description of the terms governing the use of schemas: "Perhaps there is an indication here of a preference for simple explanations—a preference that manifests itself only when the hypothesis of simple causation is not too strongly contraindicated by the data" (p. 122).

One problem with relying on behavior that violates an expectancy as a way in which to motivate people to be more "thoughtful" and systematic is simply that behavior often is not so clear-cut. Typically behaviors are complex and multifaceted, open to interpretation. It is difficult to maintain a behavioral sequence that is so concrete and clear that it prevents the possibility of negative interpretations, especially given that people are often predisposed by their prejudices to find them. In addition, this strategy of motivating people to attend carefully to others requires that the burden be placed on these others not only to articulate valid and engaging ideas, but to worry about how these valid and engaging ideas are being presented. Rather than having perceivers motivate themselves to think deeply, it asks others, the targets of these perceptual processes, to metaphorically shake perceivers from the slumber of their heuristic processing. In Allport's (1954, p. 384) terms, this state of affairs places a "preponderance of responsibility" on other people. But must the motivation to engage in systematic processing be initiated by the behavioral strategies of others? Let us examine next how perceivers can, of their own accord, be motivated to engage in systematic processing.

Raising the Sufficiency Threshold. Rather than being overpowered with clear and unexpected behavior initiated by others, which forces perceivers to process systematically, the perceivers can *adopt goals* that promote systematic processing. In this case, the "close looks" of which Bruner (1957) spoke are willfully chosen as a strategy: People are motivated to initiate elaborate processing because they desire greater certainty in their judgments than least-effort processing will allow. Instead of having a low sufficiency threshold (and desiring to exert minimum amounts of processing effort), people are described as being flexible enough to occasionally set high sufficiency thresholds that require elaborate processing. In this case, a confidence gap is created because *perceivers' desired level of confidence drifts upward*. When perceivers have set a high sufficiency threshold and desire a great deal of confidence in their judgments, then it is likely that the judgments produced by heuristic processing will not be experienced as good enough, and a confidence gap will exist. It is the perceivers' own goal to be highly confi-

dent in their judgments, evaluations, and actions that instigates systematic processing (see Figure 5.6).

Thus, in instances where people do not desire extreme accuracy, minimal processing effort can deliver judgments that surpass the sufficiency threshold (as seen at Time 1 in Figure 5.6). However, when the judgment matters, they are then motivated to set higher sufficiency thresholds and demand greater judgmental confidence. When desired confidence drifts upward, it may be set at such a high point that current confidence in a heuristic-based judgment falls below the sufficiency threshold and is no longer experienced as “good enough” (Time 2 in Figure 5.6). The higher sufficiency threshold will create a confidence gap if heuristic-based judgments are no longer sufficient. This gap will increase processing effort and require perceivers to engage in systematic processing to deliver a sufficiently good judgment (Time 3 in Figure 5.6). The natural question that follows is this: What goals are capable of producing this upward drift in sufficiency thresholds?

One such goal that leads to increased sufficiency thresholds is the goal of having a vested interest or increased *personal involvement* in the outcome of the judgment. For example, egalitarian individuals might set a high sufficiency threshold and demand to have accurate and confident judgments when interacting with members of socially stigmatized groups, because this situation is personally relevant to them—it addresses an important goal for them. The role of personal relevance in the use of heuristic, top-of-the-head, effortless processing was examined by Borgida and Howard-Pitney (1983), focusing on this motivation’s impact on the effects of perceptual salience. Salient things

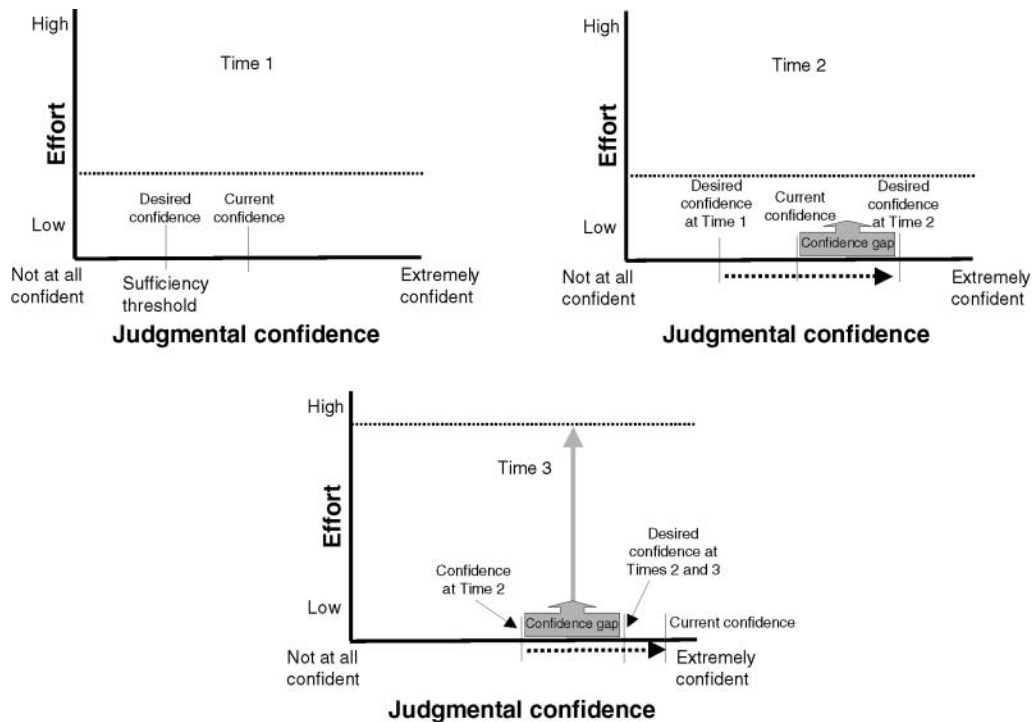


FIGURE 5.6. Instigating systematic processing by desiring greater judgmental confidence.

attract our attention. As reviewed in Chapter 1, we assume that the things on which we are focused are more likely to be the causes of what goes on around us (Taylor, Crocker, Fiske, Sprinzen, & Winkler, 1979), and we accord more weight (and perhaps more ability to influence us) to the people on whom we focus attention. The use of such heuristic and effortless processing leads to effects such as our being more likely to agree with people who sit at a head of a table, and to assign greater responsibility to the actions of people who stand out in a crowd (even if they stand out for reasons that have nothing to do with the behavior we are judging). But do these effects of attention and salience slip away when personal relevance is involved?

Borgida and Howard-Pitney (1983) had perceivers listen to a debate on a topic that they believed was either relevant or not relevant to them personally. Specifically, personal involvement was manipulated by varying whether the participants (all college students) believed they would be personally affected by a proposed change in the undergraduate course requirements. Their attention was focused on one person in the debate whom they heard arguing for one side of the issue. When personal involvement was low, students rated the salient person more favorably, regardless of whether that person argued in favor of or against the proposed changes. Perceivers seemed to be using a heuristic that dictated giving preferred status to salient people. When involvement was high, the ratings of the salient person now depended on what position the person took regarding the curriculum changes, rather than simply where the students were focusing their attention. Thus, when perceivers had heightened motivational concerns regarding a topic (i.e., a vested self-interest in it), salience effects were attenuated. Raising the sufficiency threshold so that the perceivers desired greater confidence in their judgment shifted their processing strategy from top-of-the-head processing to more systematic processing.

Similarly, individuals with the goal of not appearing foolish in public might be motivated to set a high threshold for their desired judgmental confidence. This type of goal, where people are motivated by public scrutiny and evaluation of what they think, say, and do, is called *accountability*. For example, if you are reading about a topic for a public presentation that you need to give, you may attend more closely to the material and be more careful in the conclusions you reach than if you were reading the same material for private consumption. Or, if a supervisor evaluates your performance, the amount of time and effort you put into the product you put forward is likely to be greater than if you were not being held accountable for your work. It should be noted that accountability does not always lead to more accurate thinking, but it leads to the kind of thinking desired by the audience to which the cognizer is accountable. Thus accountability could lead to self-serving thinking, outcome-directed thinking, or the like. But often, being accountable means that others are observing cognizers to see how accurate they are. Tetlock (1983) examined the impact on judgment of being held accountable by assessing the degree to which people engage in heuristic versus systematic processing in situations where they are held accountable versus not accountable.

Tetlock (1983) focused on a particular form of effortless processing—the tendency to maintain an already established belief, even in the face of evidence that might suggest revising that belief. In Tetlock's research, the focus was on *primacy effects*. As reviewed in Chapter 1, this is the finding that early information has a greater impact on judgment than later information. Thus beliefs about a person tend to be formed quickly from the initial information received, and later information is not given the appropriate weight, effort, and evaluation required to incorporate it fully into the final judgment. The key assumption of the study was that accountable participants would be less suscep-

tible to primacy effects because of how they encoded information. Accountable, as opposed to unaccountable, participants should be more cautious about drawing inferences from incomplete evidence. This should lead them to attend to more information and analyze that information more carefully. Because of such vigilance in processing, they should be able to recall more of the information at a later point in time. Consistent with these predictions, unmotivated people exhibited primacy effects in forming impressions, placing greater emphasis on information presented early. However, people who were told prior to receiving any information that they would be accountable for their impressions were immune to being influenced by the order in which that information was presented. In addition, measures of how much information was recalled revealed that accountability led to better memory for information.

One final example of a goal that leads people to raise their sufficiency thresholds is the goal of being accurate. Sometimes the situation explicitly calls for this goal, such as when people serve as jurors (where a judge might instruct them to pay close attention and to be objective and accurate) or work as teachers (where evaluating students requires accurate assessment of exams). Other situations merely suggest (imply) that accuracy is important, such as working with other people to attain some mutual goal (e.g., you have a partner for a joint project, and being rewarded depends on the performance of your partner). When one person is dependent on another individual to obtain a desired outcome, and that outcome cannot be attained through the person's own effort alone, the person is said to be in a state of *outcome dependency*. In the example just given, it is useful to form accurate judgments about what your partner is able to contribute, so you can best utilize those skills; you need to assess him/her accurately.

Neuberg and Fiske (1987, Experiment 2) examined how outcome dependency alters people's reliance on heuristic, categorical, and effortless processing. They reasoned that when a negative expectancy is associated with a particular category, people typically evaluate members of that category by using the negative expectancy as a guide. For example, most people hold negative expectancies about what it would be like to interact with a person who has schizophrenia. Thus they would be likely to form judgments of a schizophrenic person that are fairly negative, despite not thinking too carefully about or attending much to the behavior of that person. The person is treated not as an individual, but as an instance of the category. However, individuating and systematic processing could perhaps break this habit of categorical processing if the perceivers had sufficient motivation to perform this effortful type of processing (if their sufficiency threshold drifted upward). Neuberg and Fiske asked participants in their research to volunteer for a purported program on "patient reintegration" in which patients would be released from the hospital and need to start preparing for life on the outside. Thus nonthreatening interactions, in which the patients with schizophrenia worked on tasks with others (the participants), were said to be required. For some participants the patients were said to be schizophrenic, and thus negative expectancies were introduced via this category. Other participants had no category-related expectancies. The participants were then led to believe that they were either outcome-dependent on their partners in this interaction, or fairly independent of the partners. Participants were said to be outcome-dependent when they were told that their joint performance with their partners could win them a cash prize if their teams performed better than other teams. Finally, supposedly in an effort to help strangers work together in this fashion, the participants were told that they would be exchanging some personal information with their partners before working together on a series of tasks.

The main question of interest was whether being outcome-dependent would make the participants evaluate the information they received about others more closely. If there was no category label, then no negative expectancy would exist, and people should form impressions of their partners that were based on the details of the information they received (there would be no heuristic cue to suggest a heuristic, expectancy, or stereotype that could easily guide judgment and explain the behavior). If there was a category label, impressions of the partners should be based on relatively effortless processing, derived from the expectancy rather than from a detailed evaluation of the information provided. However, participants should not simply rely on their expectancies if they had a motivation to be accurate when judging someone for whom a category label had been provided. And having their partners linked to their outcomes should make people want to be accurate, so they could work best together and get the money. Thus the researchers predicted that if positive information was provided, participants would form a positive impression of their partners. If negative information was available, ratings would be negative—not because people were simply using their expectancy and were heuristically processing, but because they had raised the desired amount of confidence they needed in the judgment of their partners, had spent time thinking about these persons, and had obtained negative information. Therefore, it was predicted that effortless processing should only occur if the partners were identified (labeled) as being part of some group to whom negative expectancies were linked and if the partners were not outcome-dependent. These predictions were supported by the research results. The experimenters timed how long it took perceivers to form their impressions of their partners and used this measure as a proxy for the amount of processing effort. As predicted, when people were made outcome-dependent on persons labeled as having schizophrenia, they took twice as much time to form their evaluation of those persons as the people who were not outcome-dependent did.

The Theory of Lay Epistemics

As noted in the Introduction to this book, *epistemology* is the study of knowledge and knowing. The theory of *lay epistemics* (Kruglanski, 1980, 1990) is a theory of how laypeople come to know what they believe they know. According to the theory, the acquisition of social knowledge occurs through two phases or processes. In the first phase, people generate hypotheses about the reasons for (causes of) the events in their social world. In the second phase, people then test and validate these hypotheses. The first process presents an intriguing problem. A potentially infinite number of hypotheses may be generated to explain any given event; therefore, the process of generating hypotheses must somehow be stopped so that people can select, and feel confident in the selection of, one of those hypotheses. Once they have settled on a hypothesis, the persons have now arrived at social knowledge in which they feel confident and that they can use as a sufficient explanation for the event (and to plan how to act). They have moved from deliberating endlessly about the possible reasons why an event has occurred to having a satisfactory answer.

Let us consider an example. You and your colleagues/friends are trying to decide who to admit into your ranks as an employee/friend (e.g., to admit someone into your basketball team, work group, company, social club, etc.). Your colleague Angelika argues strongly against one candidate and strongly in favor of another, and you, as the lay perceiver, try to understand why she has argued as she has. The social event that

needs to be explained is Angelika's comments. A number of hypotheses may rush to mind to explain her remarks. First, maybe Angelika truly thinks one candidate is the best and another is the worst. Second, it is instead possible that she believes that most candidates will polarize the group and create too much dissension between group members (some of whom like one person, while others despise that person); thus Angelika has chosen to advocate for the person she believes everyone will agree is acceptable, rather than the person she really thinks is the best. Third, perhaps Angelika is aware that the ultimate responsibility for the decision lies in you, and wants nothing more than to help make you look good (and has chosen candidates that will bring prestige and honor to the group and to you). Fourth, perhaps she knows that responsibility for the decision lies in you and wants nothing more than to humiliate and embarrass you, so she has tried to manipulate the situation so as to produce the worst possible solution. Fifth, maybe Angelika has a personal agenda. She knows that such decisions are complex, and wants to admit someone to the group who helps the group on some dimension that she feels is important, even if the criterion she is using to make her decision is not one valued by you or the rest of the group (e.g., most group members want someone who is highly creative, but Angelika disregards that criterion and instead wants someone mathematically inclined and technologically savvy, regardless of the candidate's creativity). Finally, perhaps social pressures external to the actual decision are exerting an impact, such as affirmative action policies (e.g., "We must hire a Black person") or prejudice ("We must not hire a Black person"). As this example demonstrates, almost any social event poses many possible causes. As perceivers we begin immediately, and somewhat automatically, to generate such explanations. The theory of lay epistemics focuses on when and why and how people shift from generating and deliberating over a multitude of hypotheses to confidently adopting a specific, single explanation.

Thus, in order for a person to arrive at a judgment, the epistemic process must be brought to a halt. In the language of Gestalt psychologists, the stream of hypothesis generation must be "frozen" so that the cognitive system attains a state of equilibrium, or *closure*. As discussed in the Introduction to this book, when people have uncertainty or doubt, a state of psychological tension associated with lacking closure is experienced. People are motivated to reduce this tension, put an end to doubt, and attain closure; they arrive at a sufficient explanation for why an event occurred. Thus, while generating many hypotheses regarding why an event occurred can be perceived as a good thing, it also introduces an extension of a period of uncertainty, doubt, and lack of closure. At some point, this period of uncertainty must be brought to an end and closure attained. Kruglanski's (1990) theory of lay epistemics focuses on the factors that determine when people attain closure—that is, when they stop generating hypotheses and decide they have sufficient confidence in an hypothesis that has already been generated and validated.

The theory of lay epistemics outlines two general categories of factors that determine when hypothesis generation will stop. One factor is related to *cognitive capability*. People's capability to generate hypotheses involves their knowledge of the person/topic/event at hand, as well as momentary factors in the situation that promote one hypothesis over another. For example, one may not be able to generate multiple hypotheses about an event because one knows relatively little about the event. One's capability to generate hypotheses is not determined solely by what information is or is not available. The situation one is in may also limit the number of hypotheses one is capable of generating by making some hypotheses more salient than others. In the

running example, if something in the current situation reminds you how selfish Angelika can be, you may generate only hypotheses that are related to her being selfish or opposed to you doing well, and you may lack the capability to generate hypotheses that would allow you to consider her as trying to be fair or to help you do well. The situation can restrict what comes to mind by making hypotheses that are related to the current situation more likely to be generated, while unrelated hypotheses do not come to mind.

A second factor that determines when hypothesis generation will come to a halt is perfectly unrelated to whether people are capable of generating hypotheses; it is the *motivation to engage in hypothesis generation*. People's goals dictate how willing they may be to suspend having closure while considering all the possible alternatives and generating hypothesis after hypothesis. Certain motives lead people to deliberate effortfully about an event, with many possible hypotheses being generated and evaluated. At other times, people are simply not motivated to exert the effort of generating hypothesis after hypothesis and evaluating each. Instead, they are motivated to stop thinking fairly early in the hypothesis generation process and to arrive at a fast solution, with relatively little cognitive effort being exerted. Thus goals serve to energize and deenergize the epistemic process. What sorts of goals does the theory of lay epistemics suggest control the type of thinking in which people are willing to engage? Two broad classes of goals determine whether the epistemic process will be frozen early after consideration of only a few hypotheses, or thawed over a longer period of time after a systematic examination of many hypotheses. The theory describes these as goals that promote a *need for structure* (in later versions of the theory, this terminology has been changed to a *need for closure*) and goals that promote a *desire for validity* (or a *fear of invalidity*). A need for structure/closure will promote attempts to attain closure quickly, while a desire for validity/fear of invalidity will lead people to suspend attaining closure until the best possible conclusion is reached.

Need for Structure/Closure

Need for structure is characterized by quick cessation of the epistemic process; it elevates the costs associated with lacking cognitive closure, leading people to seek the quickest route to closure. It instigates a search for "an answer on a given topic, any answer, as compared to confusion or ambiguity" (Kruglanski, 1990, p. 337). It promotes a state where alternative hypotheses are less likely to be generated, information inconsistent with what is expected is less likely to be attended to, and ambiguity is avoided (making it conceptually similar to Frenkel-Brunswik's [1949] notion of intolerance of perceptual ambiguity, which describes people as thinking in a manner marked by "perceptual rigidity, inability to change set, and tendencies to primitive and rigid structuring," p. 122). Neuberg, Judice, and West (1997) make an important conceptual distinction between different forms of closure seeking—*seizing* and *freezing* processes. They assert that there are two largely independent epistemic motives relating to halting the epistemic process. The first is called *seizing*, and is determined by the desire for what is called *nonspecific closure* ("Any decision will do"). By this it is meant that perceivers are motivated to have an answer quickly, but it can be any answer that is reasonable and makes sense. They seek closure, but it need not be one specific type of closure. Independent to this is a process called *freezing*, where the focus is on attaining one form of specific closure—the desire to make a decision consistent with an existing (even if recently obtained) structure.

If a heightened goal to attain structure leads people to arrive at quick and easy answers (grabbing at the first reasonable explanation for a given social event), Kruglanski and Freund (1983) reasoned that people's judgment would be more likely to be influenced by expectancies or inferences that were formed early in a judgment task. Such information should have a greater impact on people with a heightened need for structure, because their lack of generating and examining alternative hypotheses makes it less likely that their initial judgments will be modified in light of other information that might be later encountered. Thus the impressions of persons with a high need for structure should be more predictable from the evaluative implications of traits presented early in a sequence than from the evaluative implications of traits presented later (a primary effect). To manipulate need for structure, Kruglanski and Freund placed some of their research participants under "time pressure" by asking them to make judgments within an extremely short time frame. Kruglanski (1990, p. 337) stated that an important cost of lacking fast structure is "failing to act or decide on an issue in time to meet an important deadline." Under time pressure, the threat of missing such deadlines presumably looms large and is assumed to heighten need for structure. After half of the participants were placed under time pressure and half were not placed under time pressure, they were presented with pieces of information about another person and asked to rate that person's qualifications as a job candidate. Presumably, if participants were playing the role of an employer evaluating a job candidate, they would be motivated to form an accurate and valid judgment and to consider all the information being presented equally well. But if need for structure promoted fast closure and settling on the first satisfactory conclusion, then the information presented earlier in the list of information provided should be more influential in the ratings of the participants. To test this prediction, half of the participants received positive information about the candidate followed by negative information, while the other half received negative information followed by positive. As predicted by the theory, participants exhibited a higher degree of primacy effects when there was a high need for structure: People under time pressure who saw negative information first formed more negative impressions than people who saw negative information first who were not facing a deadline. A similar pattern emerged when positive information was presented first: More positive impressions were formed when people had high need for structure.

Desire for Validity (Fear of Invalidity)

A danger of reaching a conclusion too quickly (such as when one has a high need for structure) is that the judgment may not be a valid one. *Fear of invalidity* is the fear of committing oneself to an invalid position (especially if negative consequences are associated with invalid judgments). The processing initiated by a fear of invalidity is characterized by an epistemic thaw (avoidance of early closure) and an extension of the epistemic sequence, so that one has a prolonged examination and extensive processing of relevant information. Fear of invalidity can motivate people to be energy users—cognitive philanthropists rather than misers. It can be instantiated by a desire to be accurate, by the association of punishment with being inaccurate, and by a desire to avoid the dullness and predictability associated with relying on whatever thoughts come to mind first. Mayseless and Kruglanski (1987) manipulated desire for validity by telling some research participants that the task they were performing was one that had "considerable functional significance, and constitutes a component of general intelligence" (p. 176). Essentially, they were led to believe that reaching a decision too quickly on the

task, without examining the relevant information extensively, could reveal that they lacked intelligence. Thus, with this negative consequence associated with forming an invalid judgment, participants were assumed to be motivated to avoid reaching closure too quickly. Other participants were told nothing about how the task related to their intelligence. The task that they were asked to perform involved looking at enlarged pictures of common objects that were taken from unusual angles. This made the objects difficult to identify. Participants were asked to list all the conceivable hypotheses they could generate regarding what type of object represented in a given picture. Relative to the control condition, participants who feared making an invalid judgment and were motivated to avoid early closure generated significantly more hypotheses.

Thus the theory of lay epistemics describes two processes of hypothesis generation and cessation that people use in attempting to reduce doubt. The theory posits that individuals perceive a stimulus, generate hypotheses about what the appropriate category to “capture” that stimulus might be, test (validate) the hypotheses, freeze the hypothesis testing when an appropriate solution has been arrived at, and then make inferences and plan action accordingly. The particular course (i.e., the evidence attended to and the weight given a piece of evidence) and termination of the validation process are directed by whether the persons desire an answer quickly (their willingness to “live with” uncertainty).

Correction or Dissociation Models

In the dual-process models reviewed thus far, we have seen one type of effortful mental process that people employ: They think more deeply, reassess prior beliefs, and analyze features and details more closely, engaging in counterargument generation and attempts to validate, through close inspection of the veracity of the information at hand, the data before them (whether these data are the qualities and behaviors of another person, the issues and arguments heard and observed in a commercial or political speech, or their own prior beliefs). We turn now to another type of mental task that requires effortful/systematic processing—correction processes.

Some dual-process models share the assumption that people are sensitive to and nervous about biases that may be influencing their judgments and impressions of others. People like to think that their impressions are fair, sufficient, and accurate, regardless of whether this is actually the case. What truly matters to them is their ability to believe firmly (with confidence) in the appropriateness of their thoughts. When they become suspicious that an unwanted bias is affecting their thoughts, it triggers an attempt to correct their judgment and remove this unwanted influence. The desire to have an uncontaminated and appropriate set of cognitive processes initiates mental effort bent on delivering this “purity” of thought. Thus the motivation to be accurate and fair, to have uncontaminated and unbiased judgments and decisions, will lead people to exert increasing amounts of mental energy and effortful processing.

Typically, these types of models assume that by default, people engage in a more effortless, categorical type of information processing. Such processing, as we have been exploring, is susceptible to many types of bias. People rely overly on heuristics; they overgeneralize from categories to individual members of those categories (often without sufficient justification); they disregard information incompatible with their expectations and place undue judgmental weight on information that confirms what they already believe; and they allow person types and salient exemplars to limit how a novel person is evaluated. So long as people remain blissfully unaware of these contaminating

forces, they are satisfied with their judgments. In addition, even if they were aware of these category-based processes, they might still be satisfied with their judgments so long as the categories were felt to be warranted, useful, and accurate enough. However, when people (1) become aware of these effortless processes and (2) also begin to suspect that the processes are exerting an unwanted influence on judgment, it is then that they feel the motivation to increase processing effort to rid their impressions of these intrusions.

In essence, an initial, category-based judgment anchors an impression, and a subsequent, effortful process of correction/dissociation adjusts that impression away from the anchor and toward a more “objective” accounting of the other person. Several models in social psychology adopt this “anchor and adjust” processing logic in accounting for social cognition, and subsequent chapters review several of these. Chapter 7 examines the dual-process models of Gilbert (1989) and Trope (1986a), which account for how and when people rid themselves of an attributional bias arising from effortless inferences about the traits and qualities of those they perceive. Chapter 10 examines the “set–reset” model of L. L. Martin (1986), which accounts for a bias in impression formation linked to the overuse of concepts and exemplars that one has recently been exposed to (information that is perceptually ready). Chapters 11 and 12 present a model of stereotyping suggesting that stereotypes are effortlessly triggered and then consciously removed from evaluations and judgments (Devine, 1989).

These models share a set of assumptions about correction processes that also reveal the delicate nature of such processes. Debiasing one’s responses, so that they are corrected for an unwanted influence, depends on one’s possessing (1) awareness that one is biased, (2) motivation to remove this influence, (3) theories about what that influence is and how it pushes judgment, and (4) cognitive resources for carrying out these effortful processes. Because of these contingencies, one may have an imperfect ability or desire to engage in the effortful processing that allows judgments to move beyond category-based and heuristic processing.

Awareness

Even when people desire to be accurate in judging others and to have no unwanted influences from their effortless inferences about others, they may simply fail to recognize that such influences exist. Defeating bias by being systematic, effortful, and deliberate requires the knowledge that a potential source of bias exists! Chapter 1 has been dedicated to expounding the idea of naive realism—a phenomenon that, by definition, tells us that people simply are not aware of the fact that they are biased, clinging instead to the false belief that their perceptions are based on some “absolute truth.” This is particularly troubling in the domain of stereotyping, because people often do not know that they use stereotypes. For instance, they think the problem of racism is caused by some troublesome group of racists “out there” somewhere, not by them. This naivete in typical persons’ recognition of their own biases was described in an address entitled “In Search of a Majority” delivered by James Baldwin to the Kalamazoo College class of 1960 (and later printed in the 1961/1993 collection of essays called *Nobody Knows My Name*):

Presumably the society in which we live is an expression—in some way—of the majority will. But it is not so easy to locate this majority. The moment one attempts to define this majority one is faced with several conundrums. Majority is not an expression of numbers, of numeri-

cal strength, for example. You may far outnumber your opposition and not be able to impose your will on them or even to modify the rigour with which they impose their will on you, i.e., the Negroes in South Africa or in some counties, some sections, of the American South. (p. 127)

I want to suggest this: that the majority for which everyone is seeking which must reassess and release us from our past and deal with the present and create standards worthy of what a man may be—this majority is you. (p. 137)

Motivation to Correct

For correction processes to be engaged, people must be motivated to stop the bias from exerting an influence once they become aware of the bias. This may seem obvious, but there are many instances where people may simply not care that they are biased (or at least care enough to do anything about it). For example, prejudiced persons are unlikely to stop using stereotypes once they become aware that their stereotypes may be affecting them. Partisans are unlikely to change how they evaluate the members of an opposing political party if made to realize that party membership may be biasing their views. A parent is unlikely to stop giving the benefit of the doubt to his/her child, even if there are many signs that doubts should exist. Thus, in addition to being aware of bias, perceivers must, through goals linked to things like outcome dependency and accountability, be motivated to remove the bias.

Naive Theories of Bias

If one is aware of a bias and motivated to attempt to remove that bias (correct for it), one still requires some set of beliefs concerning what the nature of the bias is and how to best eliminate it. One must have some theory of the manner in which those biases exert an impact. Once a theory has been identified, one can then attempt to implement an appropriate correction strategy. Theory-guided corrections work in a “do-the-opposite” manner—identifying the direction of the bias and attempting to shift judgment in the opposite direction (thus restoring judgment to “normal” and undoing the impact of the theory). Consider the following snippet from an article by Elisabeth Bumiller about President Bush’s reasons for going to war in Iraq, which appeared in *The New York Times* on September 16, 2002:

Last week the president drove the speculation even further when he referred obliquely at the U.N. to an assassination attempt against his father, saying, “In 1993, Iraq attempted to assassinate the emir of Kuwait and a former American president.” A senior official later told reporters, “Obviously, one doesn’t want to appear to personalize this.” In fact, only two men know for sure how much the father is influencing the son on Iraq. But interviews with Bush family friends and senior White House advisers produce a consistent belief: not much. The 41st president would never so directly tell the 43rd what to do, they say. They also insist that President Bush is able to distinguish between his feelings as a son and his responsibilities as commander in chief. “I promise you, he has the ability as a human being to disassociate any personal emotions with the job he is doing as president,” said Ron Kaufman, a Bush family friend and a close adviser to the 41st president. (p. 15)

This quote reveals several assumptions (by a variety of people) about what happens when people suspect that they are biased, or that someone or something in the environment is

biasing their responses. The first assumption is that people “know for sure” what influences them. Others may not know, but the persons being influenced can detect bias and can accurately deduce the effects of that bias. The second assumption is that a person who becomes aware of bias has the ability, “as a human being,” to dissociate that bias from responses. Thus if President Bush had personal feelings that drove him to desire revenge for his father’s attempted murder, he could detect those feelings and accurately know how to adjust his reaction so that his decision about whether to attack Iraq (the nation responsible for the alleged murder attempt) was not influenced by these personal feelings. Are such assumptions accurate? What actually happens when people think they are being influenced by some unwanted agent (either internal or in the environment)?

According to Wegener and Petty (1995), the answer to this question is that people try to correct for this bias, to prevent themselves from being influenced. This is precisely what the quote above describes. How can they accomplish this? People hold *naive theories* about the nature of the influence (theories about the magnitude and direction of some biasing force) and take steps to correct for that perceived influence by countering the impact that these theories claim exists: “People might use a subjective reaction to a target but adjust that initial reaction in light of a theory” (Wegener & Petty, 1995, p. 39). Thus, in the quote above, this is reflected in the fact that President Bush was said to “distinguish between his feelings as a son and his responsibilities as commander in chief” and to dissociate these personal emotions from his professional decisions. The subjective reaction of the President might be rage at Iraq’s leader and desire for revenge, but he could detect this influence, had a theory about how it was influencing his decision whether to initiate a war, and could remove that influence (adjust the subjective reaction to account for the detected bias). Wegener, Dunn, and Tokusato (2001, pp. 281–282) have summarized this succinctly: “[The perceiver] will evaluate the potential biasing effect(s) of salient factors in the judgment setting. . . . This is accomplished by consulting naive theories of the bias(es) associated with the salient factor(s).” However, just because people suspect they are being influenced and take steps to try to prevent it, this does not necessarily mean that they are correct in assuming an influence exists or that they are taking the appropriate steps to fix it if it really does exist.

For example, a teacher, Mr. Vance, who has his own child in class might worry that his positive affect for the child could be influencing his evaluation of the child’s ability as a student. When concerned with contamination of this sort in one’s subjective reaction, one attempts to correct it. One holds a naïve theory about that influence, and this theory then shapes the type of correction that is attempted. In this example, Mr. Vance has a theory that he is being too lenient toward his own child, so he acts in an opposite fashion: He starts being extra critical of the child and using harsher standards than those he uses for other students. However, these are referred to as *naive theories*, because such theories of influence need not be accurate. What triggers the use of theories of influence is the *perception that one is biased or is being unduly influenced*, not an actual influence. Even if Mr. Vance is not being biased in favor of his own child, he could still believe this bias to exist and still hold a theory about the nature of the bias and how to rectify it. Nisbett and Wilson (1977) describe an experiment that provides an excellent illustration of how people perceive an influence on their mental processes where none really exists. Research participants watched a film either with or without a distracting background noise. When participants in the background-noise condition were asked whether the noise influenced their ratings, most reported that it had. However, the actual ratings of the film did not differ for these two groups. Their theory that they had been influenced by the noise was simply incorrect.

One may not only be using a theory of influence when one is not really being influenced; one may end up using the wrong theory (e.g., Wilson & Brekke, 1994). Thus one may really be biased, but have the wrong naive theory about how strongly one is biased (the magnitude of bias) or in what way one is biased (the direction of bias). In our example, Mr. Vance may truly be biased toward his child, but only in that he has a moderately positive predisposition. If he then acts in an extremely harsh way to compensate for the bias, his theory has overestimated the magnitude of the bias. Alternatively, he could have a flat-out incorrect bias: He could actually be more harsh toward his child, but he may have a naive theory that he is probably too lenient on the child. He may then use the theory to trigger even harsher treatment. Bias does indeed exist, but the naive theory Mr. Vance uses to inform him of the nature of the bias is wrong (or in the wrong direction). In the President Bush example, the intuitions expressed in *The New York Times* about how the President might deal with his emotions toward Iraq were partially correct. As a human, he would try to detect the bias and correct for it. However, the assumption that he would detect the correct bias and accurately dissociate this bias from judgments he made as commander in chief was not a good one. Let us now provide a few empirical examples of how people use naive theories to try to correct their judgment, and as a result produce what may be bad judgments.

Wegener and Petty (1995) assert that people hold naive theories reflecting an idiosyncratic view of how they are being influenced. If true, then the manner in which people attempt to correct their judgments should reflect the underlying theory of how they are being influenced. If one has a theory that one is being influenced by having one's judgments pushed in a positive direction (such as Mr. Vance's favoring his own child), then one should try to make corrected judgments that are in a more negative direction (such as Mr. Vance's judging his own child more harshly). If one's naive theory is that judgments are perceived to be influenced in a negative direction, then one tries to be more positive to counteract the perceived bias. Wegener and Petty use a more precise language, distinguishing between *assimilative* and *contrastive* biases. In this terminology, there is a context, and one has a naive theory that one is being influenced by this context in one's judgment of some target. If the target is made to seem similar to the context as a result of this influence, this is said to be an *assimilative influence*. If the target is made to seem more unlike the context than it would had the context not been present, this is a *contrastive influence*.

For example, Wegener and Petty (1995) asked participants how they thought that thinking about the weather in exotic locations like Jamaica and Hawaii would influence their ratings of less exotic places (Midwestern cities like Indianapolis) as vacation spots. They predicted that most people would believe it to have a negative impact, or a contrastive influence (they would have a theory specifying that less exotic places will seem extra undesirable in comparison to Hawaii). They also asked participants how desirable the jobs were perceived to be in beautiful weather locales, like Hawaii. In this case, an assimilative influence was predicted such that because of the great weather, people would believe that others probably love working and living there. A first experiment revealed that these were precisely the naive theories that research participants held—the same context (thinking about the weather in Hawaii) led to a perceived contrastive influence on one judgment target (Indianapolis) and an assimilative influence in another (working in Hawaii). Once having identified these theories, the researchers could now explore the real question of interest to us: If people have different theories about how they are influenced when judging different targets, do they correct for the influence, and do the corrections taken differ when judging each of the tar-

gets (according to the theory of influence specific to that target)? Wegener and Petty expected that the negative impact on ratings of Midwestern cities from having contemplated the exotic locales would be corrected for. These cities should increase in attractiveness if people were aware of the bias and were using naive theories to correct for it (recognizing the impact of the biasing agent and removing that influence from judgment). To assess this, participants were asked to rate the desirability of vacationing in Midwestern cities, and to be sure not to let their ratings of exotic locales influence them (thus making sure that they would use their naive theories). The results revealed that the more negative the theories of influence (such as theorizing one would harshly judge a Midwestern city after thinking about Hawaii), the more positive subsequent ratings of these Midwestern cities became. People adjusted their judgments, using their theories of influence as a guide.

Wegener, Petty, and Dunn (1998) illustrated this point further, using people as the targets (rather than cities). For example, research participants had to rate how violent Arnold Schwarzenegger was. Some of the participants made this rating after first making a similar rating about extremely nonviolent people, such as Gandhi and the Pope. Other participants made this rating after first making a similar rating about extremely violent people, such as Hitler and Stalin. The logic was that the initial ratings of either extremely nonviolent or violent people would serve as a context that could introduce bias. Perceivers would detect this potential source of bias and worry that their ratings of Schwarzenegger would be influenced by the previous ratings. One group would hold a theory specifying something like "Next to the Pope, Arnold is pretty violent," whereas the other group would hold a theory with exactly the opposite implications: "Next to Hitler, Arnold is not violent." Thus the same target (now Governor Schwarzenegger of California) would be reacted to in extremely different ways if people followed their theories. Participants who first rated Gandhi would worry that they would be biased by Gandhi's nonviolent shadow and rate Arnold as extra violent. Thus, to be nonbiased, they would have to make sure they rated Arnold as nonviolent. Participants who first rated Hitler would worry that they were being biased by Hitler's genocidal nature and would underestimate Arnold's violence as a result. Thus, to be nonbiased, these people would have to rate Arnold as extremely violent. This was precisely what was found. When asked to try not to be biased and to correct for any unwanted influence on their evaluations, people holding opposing theories about how they might be influenced judged the same person in diametrically opposing ways. Naive theories of influence exist, and they guide judgment. With such naive theories in hand, and the motivation to use them, people can now correct. Or can they?

Cognitive Capacity

Even if people are aware of bias, motivated to eliminate that bias, and in possession of a theory about how (and how much) that bias is influencing them, they may still end up not engaging in correction processes. Why? Because such processes are effortful and require the systematic reevaluation of an already existing impression. As we have discussed in several places in this book, such processes require processing capacity to be available. If cognitive load is present in the current situation, the requisite capacity to engage in theory-based correction may simply be unavailable. Thus people may have the desire to correct, but lack the cognitive capability to do it. Bruner (1957) summarized this irony of social life quite nicely: "It is doubtful whether a therapeutic regimen of 'close looking' will aid the misperceiver much. . . . [The] greatest difficulty rests in

the fact that the cost of close looks is generally too high under the conditions of speed, risk, and limited capacity imposed upon organisms by the environment” (pp. 141–142). By *close looks*, Bruner refers to attempts by perceivers to analyze the data closely. Such attempts are blocked by the ever-present perils of social life—the need to think and act quickly (speed) and the bombardment of information that must simultaneously be processed (limited capacity). The processes required for correction are simply too strenuous, given the other, simultaneous processing demands. In fact, the lack of such abilities is one reason we have claimed that categorical and other forms of effortless processing have developed in the first place.

The Timing of the Goal

The discussion of how goals affect cognitive processing in the dual-process models reviewed above must also include an examination of the question of the timing of the goal’s instantiation. Tetlock and Kim (1987) suggested that goals are most successful in defeating passive influences on judgment by making people more vigilant in the way they *encode* information. That is, if you are going to exert effort in thinking about a person, it is best to do so at the moment that you are receiving that information and actually interacting with that person. This is not to say that goals cannot have an impact on your *reprocessing* of information, or on your attempts to retrieve information about a person from memory and engage in systematic processing on this stored information. However, the systematic processing of information from memory requires you to have stored enough information about that person in memory that you can later retrieve it and process it elaborately (Thompson, Roman, Moskowitz, Chaiken, & Bargh, 1994). Obviously, the inclusion of this extra step will mean that systematic reprocessing is at a relative disadvantage to systematic processing that occurs at encoding (online).

For some types of categorical/heuristic/schematic processing, specific influences on information processing can be demonstrated by varying when goals are provided. For example, in the research of Tetlock (1983) reviewed above, accountability goals triggered systematic processing that was able to overturn judgments that were more effortlessly formed. However, Tetlock also found that this effect was restricted to instances in which people were told they were accountable *before* receiving information about the person. When they received the information first, and only later found out that they were accountable for their impressions, categorical impressions persisted (people were still heuristically processing).

However, other research has found that “when” the goal is introduced is not an important factor. Accountability triggers systematic processing when it is introduced before perceivers meet a person, as well as when it is introduced at some point after perceivers have met the person and have already formed an impression (e.g., Thompson et al., 1994). Goals (such as accountability) introduced after initial inferences have been drawn should have a *harder* time debiasing or correcting inferences, but this does not rule out debiasing altogether. If this is the case, why then does some research suggest that correction is only possible when the goal is presented in time for online correction, allowing people to generate complex and multifaceted inferences that can defeat passive influences as the information is being interpreted? Some speculation is offered.

First, perhaps some types of categorical processing are weaker than others and are more easily overturned, even after an impression has already been formed. Second, perhaps the nature of the goals needs to be particularly self-involving in order for repro-

cessing to be successfully initiated. Third, perhaps this has to do with judgmental confidence. Judgments that result from some types of categorical processing may be held more confidently than those resulting from other types of categorical processing. For these confidently held judgments, people may feel so confident that they do not feel as if they need to go back and reexamine the information. Fourth, perhaps the time itself is the key issue. If the goal is introduced after people have already formed an impression, but the time between having formed an impression and being told that they will need to be accurate is short, people will still have the ability to recreate the original information from memory and alter their impressions accordingly. But if the time interval is longer, it is possible that people's ability to recall is interfered with, so that the information is simply not available in memory any more. As previously stated, capacity alone may not be enough to instigate control over passive processes, but neither may be having intent. Regardless of what people intend to do, they must have the capacity and the ability to carry out those intentions. Perhaps these limiting conditions for correction (a goal, capacity, and ability) are more constraining at recall and make retrieval-based influences from intent more difficult to produce.

EFFORT AND ITS RELATIONSHIP TO ACCURACY AND CONFIDENCE

Biased Systematic Processing

Increased effort in thinking about a person, issue, or event is not necessarily the same thing as more accurate thinking—a fact labeled by Chaiken and colleagues (1989) as *biased systematic processing*. Systematic processing can be biased for at least two reasons. First, we perceivers may be motivated to come to a particular conclusion, expend effort to defend that conclusion, and make sure it is the result that is attained. Kunda (1987) has referred to this as *motivated reasoning*. Second, we may simply try to be accurate, but have false knowledge and incorrect theories that we use in our detailed analysis. Try as we might to be fair, if we are armed with bad information, we will end up with a false conclusion. This is called *misdirected decontamination*.

Misdirected Decontamination

Even if we become aware of a bias in our judgment, and we have the desire and ability to correct for this bias and “decontaminate” our thoughts, we may have an incorrect naive theory of what that influence is (e.g., Nisbett & Wilson, 1977; Wilson & Brekke, 1994). Thus, rather than “correcting” judgment, we adjust it (according to this wrong theory) in a manner that makes it biased in a new and different way. This is discussed further in Chapter 10. For now, let us simply highlight the fact that greater effort in thinking about people and events does not guarantee that we will have greater accuracy in our final judgments, decisions, and evaluations. One reason for this is that we service our cognition with incorrect theories about what is influencing us. In trying to fix our cognition, we pull a faulty tool out of the cognitive toolbox. For example, President Bush relied on faulty intelligence information in deciding to invade Iraq in 2003.

As another example, in his second term as President of the United States, Bill Clinton was constantly receiving negative press: He was impeached by Congress,

accused of having an affair with a White House intern named Monica Lewinsky, found to have lied to a grand jury (and to the American public in a famous, finger-wagging denial) regarding that affair, and subjected to scrutiny regarding potential illegal activity in a real estate deal (Whitewater) and in campaign financing. Ironically, the public loved him (if we use approval ratings as an index). Many people felt that despite all these troubles, he still would have defeated George W. Bush for President in 2000 if he could have run for a third term! One possible explanation for this Clinton love may have been the use of a bad theory of bias. People may have assumed that the negative affect they felt toward Clinton resulted from the intense media scrutiny and unyielding attacks of partisan Republican foes. Rather than these attacks' being successful, they may have backfired on Republicans by allowing people to develop a false theory that their negative feelings toward the President were caused by the attacks of the media and partisan opponents. Instead of a theory that attributed their negative feelings to the President's actual misconduct, people may have used a naive theory that led their thinking about the President in a totally different direction. Note that in this example people were not thinking effortlessly or categorically. They were using a theory to help them effortfully sort out their cognition about a complex target (President Clinton). It was just that in their effort to drive toward an accurate end, they were using a bad map.

Motivated Reasoning

Not all bias is due to the "innocent" inability to hold an accurate theory of influence. Some bias is quite motivated. To assume that the only motive that drives us to think deeply about information is one that specifies a drive for accuracy is to grandly underestimate ourselves as perceivers. Chapter 1 has identified a myriad of motives and goals that provide the energy driving human social cognition. The consequence of having so many motives that may potentially drive the way in which we think about others is that the energy and effort we expend on systematic processing is not always spent on gathering the best and most objective data. Instead, we rationalize, spin, and self-deceive; we selectively attend and evaluate; and we engage in a variety of mental gymnastics aimed at using our effortful cognition to prove what we want it to prove and what we desire to believe is true. This type of motivated reasoning, capable of producing bias in impression formation, is the focus of Chapter 8, so let us limit discussion of these issues here to a few relevant examples to illustrate the point. Chaiken and colleagues (1989) explored three general types of motivated reasoning.

The first type of motive that directs systematic processing is *accuracy motivation*. It is represented by goals, already discussed, that lead people to seek to be accurate: fear of invalidity, desire for self-assessment, outcome dependency, accountability, and increased responsibility. People motivated to be accurate in evaluating their own behavior are said to have *self-assessment* motives (e.g., Trope, 1986b). When accuracy-motivated, systematic processing proceeds in a more objective and unbiased fashion in an attempt to seek the "truth." Illustrations of accuracy-driven processing have already been provided in connection with the theory of lay epistemics (e.g., Mayesless & Kruglanski, 1987) and the HSM (e.g., Maheswaran & Chaiken, 1991).

Other motives besides accuracy, however, can bias perceivers. A second class of motives detailed by Chaiken and colleagues (1989) is *impression motivation*. This is the desire to maintain beliefs, attitudes, judgments, and actions that allow others to form a specific type of impression of one in accordance with one's interpersonal needs in a

social situation. It is the motive to produce a certain premeditated set of consequences in one's interpersonal relationships and to project a certain image. This is typically accomplished by trying to take into account the qualities, views, and disposition of one's partners in a social situation. Such an impression management (e.g., Schlenker, 1980) or self-presentational (e.g., Jones, 1990) goal allows one not only to anticipate the characteristics that others wish to see, but to produce the desired outcomes from another person by displaying those desired characteristics for them. Need for approval, fear of rejection, desire for power, ingratiation, communication goals, and social role needs are all motives capable of triggering impression-motivated systematic processing. Such processing is often selective and strategic, aimed at assessing the social acceptability of beliefs and judgments.

For example, Sedikides (1990) gave some research participants the goal of communicating a particular type of impression to other people, asking them to tailor their message to suit the audience who would be receiving it. Others were given no such instruction. The perceivers who were busy attempting to match their own behavior to the qualities of the people they would be interacting with were more systematic in their evaluation of these other people. When participants had no such goals, they were shown to judge other people according to whatever concepts came most easily and readily to mind. This top-of-the-head type of processing was abandoned by people who were concerned with having others form a specific impression of them. Participants were not aware of the fact that they typically process in an effortless fashion. Thus, their failure to rely on such processing was not due to a correction process or to decreased confidence in their judgment. Instead, they increased their desired confidence, with the impression goal raising their sufficiency threshold, making them want to think more deeply.

Finally, the last class of motives discussed by Chaiken and colleagues (1989) includes those aimed at protecting and maintaining self-esteem; this class is labeled *defense motivation*. The defensively motivated systematic processor scrutinizes information in a directional, selective manner, so as to confirm the validity and prevent the falsification of important, self-relevant knowledge, beliefs, and relationships. The goal of the defense-motivated person is to preserve, verify, and enhance the self-concept and those aspects of the world in which the self is personally vested. These motives reflect a desire to form and defend conclusions that are consistent with prior knowledge (including self-knowledge), as well as to form and defend conclusions that promote self-esteem. Thus, within defense motivation are two distinct types of motives—one geared toward consistency and one toward ego defense/enhancement.

Consistency Motivation and Self-Verification. Two decades prior to the rise of dual-process models, social psychology was replete with another group of models, the *cognitive consistency* theories. These theories are far too numerous and complex to review in detail here. However, the premise of such theories is that humans are fundamentally driven to avoid inconsistency among their cognitions. Festinger's (1957) *theory of cognitive dissonance* is perhaps the best known. Its premise is this: When a person has two cognitions that are inconsistent, this arouses an aversive drive state, similar to hunger. As with physically derived drives, people are motivated to reduce the discomfort associated with the drive state (the dissonance). They are motivated to restore consistency among the discrepant cognitions to reduce their discomfort. As discussed in the Introduction, such inconsistencies are experienced as doubt, and in an attempt to reduce doubt people try to attain certainty, predictability, and control over the dissonant

cognitions. This excessively brief review of consistency theories highlights the current point: People expend mental energy not only to attain accurate information, but also to attempt to maintain existing beliefs. These attempts at maintaining consistency occur even in the face of dissonance-arousing information that is inconsistent with and potentially damaging to those beliefs. As reviewed in Chapters 3 and 4, rather than overhauling their existing belief systems, people instead engage in elaborate processes (consistency resolution processes, rationalizations, data twisting) that will eliminate the dissonance while maintaining the old views. Swann (1990) has referred to such attempts to maintain existing beliefs concerning the self as *self-verification* (for a review, see Chapter 8).

Darley and Gross (1983) and Snyder and Swann (1978) provided classic illustrations of belief preservation, consistency seeking, and expectancy verification in the manner in which people evaluate and test hypotheses. People exert effort aimed at arriving at a specific conclusion that maintains a prior belief. Darley and Gross referred to this type of verification in hypothesis testing as the *perceptual confirmation effect*. Darley and Gross reasoned that when people have a particular expectancy, hypothesis, stereotype, or prior belief triggered, it will serve as an anchor that guides subsequent information processing. This research, reviewed in detail in Chapters 3 and 11, illustrates that people subsequently gather information in a manner bent on confirming their existing hypotheses. This can involve not only seeking out information consistent with what they already believe, but, as in the case of the Darley and Gross experiment, twisting the perception of existing data that are not actually consistent with their hypotheses so that what they perceive is a set of confirming evidence—perceptual confirmation.

Self-Esteem Defense/Enhancement. Sedikides (1993) recognized the importance of the motives to assess (accuracy) and verify (consistency), but argued in support of Allport's (1937, p. 169) claim that a person's "most coveted experience is the enhancement of his self-esteem." As reviewed in the Introduction, the maintenance of a positive view of the self is a powerful motive that drives cognition. When we are thinking about our social world, we are biased toward interpreting the information and events around us in a way that bolsters self-esteem and wards off negative self-evaluations. But self-esteem enhancement is not only attained through the way in which we think about ourselves; it is also determined by how we evaluate others. The self is not an island, but a complex array of *social identities* that link an individual to various social groups (what Brewer, 1991, called the *social self*). To maintain a positive sense of self requires us not only to be biased in how we evaluate self-relevant information, but to maintain positive views of the groups to which we are linked. This fact is the essence of *social identity theory* (e.g., Tajfel & Turner, 1986), which details how esteem enhancement affects evaluations of others:

- 1) Individuals strive to maintain or enhance their self-esteem: they strive for a positive self-concept,
- 2) Social groups or categories and the membership in them are associated with positive or negative value connotations. Hence, social identity may be positive or negative according to the evaluations,
- 3) The evaluation of one's own group is determined with reference to specific other groups through social comparison. (Tajfel & Turner, 1986, p. 16)

Thus evaluations of others are biased because these others are linked to us through social/group bonds.

In support of this idea, Tesser (1988) proposed a model of self-evaluation suggesting that we can be systematically biased in how we evaluate others, ranging from those with whom we share to those with whom we do not share a relevant social bond. In this *self-evaluation maintenance model*, it is assumed that others often play the role of “comparison standards.” When we seek to evaluate ourselves—to come to understand whether we are performing well or poorly on a particular dimension—we look to others for information. For instance, if students want to evaluate their performance on an exam, their objective scores are usually not enough. In addition to knowing the grades they received on the exam, individuals want to know how well they did relative to others! We all use others as a standard of comparison against which we evaluate and judge, engaging in what Festinger (1954a) called *social comparison*, for the purpose of attaining self-knowledge. Festinger stated that “there exists, in the human organism, a drive to evaluate his opinions and his abilities. . . . To the extent that objective, non-social means are not available, people evaluate their opinions and abilities by comparison respectively with the opinions and abilities of others” (pp. 117–118).

When we engage in social comparison, who do we choose to compare ourselves against? We select people whom we expect to be similar enough to yield useful information about our abilities. If I want to know whether the present book is successful, I will not compare the sales figures against the top book on *The New York Times* best-seller list. I instead will turn to similar books written by psychologists (or other academics). Additionally, since comparison to others is done to help one feel positive about one’s self, there is another factor to consider: the direction of the comparison. We sometimes engage in *upward social comparison* with someone who is superior to ourselves, because such comparisons can set standards to strive for and may provide surprisingly positive results that can really boost our self-esteem. However, a surer strategy for boosting self-esteem is to engage in *downward social comparison*—a comparison with someone whose performance is inferior to our own. Downward comparisons will typically yield a positive view of ourselves, but if we go too low, we may be providing false feedback about our ability because we are comparing ourselves to people far worse than us. We will feel good about how well we do, but the comparison may not be informative about our true abilities.

In this way, even relative strangers provide information about ourselves, through processes of social comparison, if those others have performed in the dimension for which we desire feedback and self-evaluation. In an attempt to bolster our self-views, we can compare ourselves to others and reflect on the qualities and achievements of others in a biased way. The variety of ways in which cognitive flexibility and effort can be used to promote positive self-views and to ward off negative self-views is a focus of Chapter 8, and we will review only Tesser’s (1988) self-evaluation maintenance model at this point for the purpose of providing one illustration.

The self-evaluation maintenance model recognizes that the manner in which we compare ourselves to others depends on just how close those others are to us and just how relevant their behavior is to our own self definition. Others can range from being complete strangers to people with whom we share important social ties/bonds. The behavior of others can range from being extremely relevant to us (they have the same career) to being totally irrelevant (they achieve in a domain that has nothing to do with us). Even people with whom we share social ties—through kinship, political affiliation, friendships, social groups, ethnic groups, religious groups, alma maters, hometowns, and nationality—can vary in closeness to us and in the degree to which their behavior is relevant to us. The self-evaluation maintenance model asserts that sometimes we evalu-

ate ourselves, via the performance of others, through a process of *reflection*. We can bask in the reflected glory of the achievements of people with whom we share a social bond, even when the bond and the achievement are fairly trivial (e.g., our college football team wins a game, or someone from our hometown makes the news or achieves success in music or movies). In these examples, reflection helps to bolster self-esteem because the success of others becomes a ray of light that illuminates our self-evaluation. However, reflection is really only a useful way to boost self-esteem if the behavior of others is not overly relevant to us (is not threatening to us) and if we have some social tie to the person. If we have no aspiration to be movie stars, pop stars, or football heroes, then the success in these domains of other people can reflect on us. But it is hard to bask in the glory of someone who is succeeding in the very area in which we were hoping to succeed. In addition, reflection makes far more sense when we have some form of a social tie to the person. Reflection is not a very useful way to evaluate ourselves if the person is someone to whom we have no social tie.

Instead of reflecting on the glory of others, Tesser's model also specifies that we evaluate ourselves through social comparison. Social comparison is especially likely if the person we are comparing ourselves against is performing in a domain relevant to us and in which we desire feedback. Whereas the negative performance of others may allow comparison processes to provide a boost to self-esteem, the positive performance of others can threaten self-esteem (particularly if it is someone with whom we share a social bond). It is here where biased systematic processing can come into play. We should be motivated to see ourselves positively through social comparisons, even if the objective evidence might not warrant such a conclusion. For example, a Democrat might exert a good deal of mental energy to convince him-/herself that the Democratic candidate in a debate performed better than the Republican. The employee up for a promotion might use all his/her powers of rationalization to believe that his/her performance on the job is superior to that of other employees in line for the same promotion. In sum, social comparison processes can be used to allow us to maintain and enhance a positive view of ourselves, and this can occur through the biased systematic evaluation of and comparison to the behavior of others who are behaving in a self-relevant domain.

CODA

Dual-process models address how and when people stop being cognitive misers and start engaging in close looks. This involves discovering what factors cause people to experience doubt—a lack of confidence in their judgments. There have been two general responses to this question. One has been that this occurs when the data are so inconsistent with the prior structures stored in memory (e.g., expectancies, stereotypes, schemas) that they force people to abandon a reliance (or strict reliance) on such theory-driven processing. Unambiguous information about a person—information that is counter to expectancies and stereotypes—challenges perceivers' perspective. The ability to force prior conceptualizations onto the currently encountered person is rendered impractical. A second has been that “close looks” are willfully chosen as a strategy by flexible processors: People are motivated to initiate elaborate processing because they desire greater certainty in their judgments than least-effort processing will allow. They can desire to be accurate.

The link between motives and information processing is developed from this interplay between people's striving to exert as little mental effort as possible, yet simultaneously to hold judgments in which they feel sufficiently confident (ones that surpass a sufficiency threshold). The processing strategy that is utilized depends on where the sufficiency threshold is set and how much effort is deemed necessary for a judgment to be confidently held. There are two distinct motivational assumptions suggested by this dual-process logic. One is that motives are capable of establishing and shifting the sufficiency threshold. A variety of motives can be successfully implicated in this shifting and setting function, which serves to create gaps between actual and desired confidence. The second motivational assumption is more fundamental. It is that the existence of a gap between actual and desired confidence, the experience of feeling that a judgment is not sufficient, is in itself an aversive state that people seek to reduce. People are said to desire confidence, and lacking a sufficient experience of confidence is in itself motivating.

The dual-process notion can be seen as far back in the psychological literature as James's (1890/1950, p. 451) belief that effortless processing predominates, but more effortful processing can occur under the appropriate conditions:

The stream of our thought is like a river. On the whole, easy simple flowing predominates in it, the drift of things is with the pull of gravity, and effortless attention is the rule. But at intervals a log-jam occurs, stops the current, creates an eddy, and makes things temporarily move the other way. If a real river could feel, it would feel these eddies and set-backs as places of effort.

James (1907/1991) warned that people are neither purely theory-driven (using "principles") nor data-driven (using "facts"); rather, they possess qualities on both sides of the line:

Each of you probably knows some well-marked examples of each type, and you know what each example thinks of the example on the other side of the line. . . . Most of us have a hankering for the good things on both sides of the line. Facts are good, of course—give us lots of facts. Principles are good—give us plenty of principles. The world is indubitably one if you look at it one way, but as indubitably is it many, if you look at it in another . . . your ordinary philosophic layman never being a radical, never straightening out his system, but living vaguely in one plausible compartment of it or another to suit the temptations of successive hours. (pp. 9–10)

This is a clear exposition of the idea that the "temptations of successive hours"—fluctuations across time and situations in needs and goals—can lead people to shift from a theory-driven approach to processing the world, to a data-driven approach.