

Chapter Eleven

How the Mind Moves: Knowledge Accessibility and the Fine-tuning of the Cognitive System

Leonard L. Martin, Fritz Strack, and Diederik A. Stapel

Two monks were arguing about the temple flag as it waved in the wind. One monk believed that it was the wind that was moving. The other believed that it was the flag that was moving. Despite considerable debate, neither monk was able to convince the other of his point of view. Finally, the Master arrived, and the two monks broached the issue with him. The Master noted, "It is not the wind that moves. It is not the flag that moves. It is your mind that moves." At this, the two monks were enlightened.

Zen parable

This chapter is about the mind moving. More precisely, it explores the general proposition that social perception is not a neutral registration of objective reality, but an active construction that is influenced by concurrent processes of thought, memory, feeling, and motivation (cf. Bruner, 1992; Martin & Tesser, 1992). In elaborating upon this proposition, we discuss some of the early research (e.g. Bruner, 1957) that demonstrated the constructive nature of social perception. Then, we discuss the theoretical advancements made with regard to this issue in some of the early social cognition research (e.g. Higgins, Rholes, & Jones, 1977). Finally, we discuss some recent findings that have helped to refine our understanding of the constructive nature of social perception. The chapter ends by suggesting that the processes involved in social perception are more complex than was reflected in the earlier research. Individuals are not cognitive misers who use whatever information is on the top of their heads. Rather, they are cognitive optimizers. They have access to a variety of different types of information and they use these selectively in the service of a range of processing objectives and motivations.

A Classic Example of the Mind Moving

The 1951 football game between Princeton and Dartmouth was an especially rough one. A large number of penalties were called, and Princeton's star quarterback had to leave the game because of a broken nose and a concussion. Although students from both schools agreed that the game had been rough, they did not agree on exactly how rough it had been, and on which team had started the rough play. To gain a better understanding of this lack of agreement, Hastorf & Cantril (1954) conducted a study. First, they assessed the attitudes of some Princeton and Dartmouth students toward the game. Then, they showed these students a film of the game, and asked them some questions about the film they had just seen.

Not surprisingly, the attitudes of the students at the two schools differed. The Princeton students generally thought that the game was rough and dirty, and that the Dartmouth team had started the dirty play. The Dartmouth students also thought that the game was rough, but they were more likely than the Princeton students to see the game as fair and to see both sides as to blame for the rough play. More interestingly, though, for present purposes, was the way in which these differences in attitudes reflected themselves in the two groups' perceptions of the game film. Although students from the two schools watched the exact same film, the Princeton students saw the Dartmouth team make many more infractions than their own team, whereas the Dartmouth students saw both teams make about the same number of infractions, with their team making half the number of infractions attributed to it by the Princeton students.

In explaining these different perceptions in the face of the same objective information, Hastorf & Cantril (1954) proposed that "there is no such 'thing' as a 'game' existing 'out there' in its own right which people merely 'observe'" (ibid., p. 133). Rather, "an 'occurrence' on the football field or in any other social situation . . . becomes an 'event' only when the happening . . . reactivates learned significances already registered in what we have called a person's assumptive form-world" (ibid., p. 132). Stated differently, the objective reality of the game resulted in different subjective experiences because students from the two schools viewed the game using different previously stored knowledge structures, which, in turn, led them to attend selectively to different occurrences and to interpret the same occurrences in different ways.

This general conclusion has been supported in a variety of subsequent studies (for a review, see Martin & Tesser, 1992). What these studies have shown is that individuals do not make judgments (e.g. How much do I like my mother?) by retrieving a single invariant score from memory (e.g. my evaluation of my mother). Rather, individuals construct their judgments as needed, using previously stored information as well as information from the current context (e.g. my evaluation of my mother given that I am in a bad mood and that she just grounded me). Thus, judgment of the same stimulus by the same individual can differ depending on the context in which the judgment is rendered.

It should be noted, though, that from a constructivist perspective, these context dependent changes in judgment do not reflect an inability on the part of the perceiver to retrieve his or her *real* evaluation. Rather, context dependency is assumed to be a natural by-product of the processes by which evaluations are rendered. After all, a person who is

home alone *should* have a different interpretation of the sound of a window breaking than a person who is home with children playing in the next room. The children provide a ready explanation for the window breaking, but who is breaking the window if the person is home alone? From a constructivist perspective, this example does not reflect different interpretations of the same stimulus. It reflects different interpretations of different stimuli. This is because, from a constructivist perspective, the stimulus is not “the sound of a window breaking.” It is “the sound of a window breaking while I am home alone” or “the sound of a window breaking while children are playing in the next room.” People respond contextually. By the same reasoning, an individual who has had a heart attack *should* interpret a chest pain differently than a person who has recently eaten a bowl of extra spicy jambalaya. The stimulus is not “a chest pain,” but “a chest pain in the light of my memory of what has just happened” (i.e. heart attack versus jambalaya). The more general point is that individuals bring aspects of themselves (e.g. memories, motivations), as well as aspects of the current context, to bear in evaluating any given target stimulus. This is what is meant by saying that social judgments are constructed.

It is important to point out that in assuming that social judgments are constructed, one need not assume that such judgments are arbitrary or that there are no general rules. In fact, quite the opposite is true. There are general rules, and these rules are beginning to be understood. For example, research has shown that social judgments depend, in predictable ways, not only on the particular knowledge that individuals bring to the judgment, but also on factors such as the relation between this knowledge and the target stimulus, the timing of the activation of the knowledge, the method by which the knowledge was activated, and the perceiver’s motivational state. These are the kinds of issues we examine in this chapter. We begin by briefly reviewing some of the early research that emphasized the constructive nature of perception.

The Role of Accessibility: Some Initial Considerations

The research that is most often cited as the prime influence on the study of knowledge accessibility in social psychology today is that of Bruner (1957). The goal of this early work was to show that individuals do not respond to a direct copy of the objective world, but to their categorization of the objective world (Bruner, 1992). This is the case, according to Bruner, because most information is relatively meaningless until it has been identified with a mental category. A dark tubular object, for example, would elicit little reaction in an individual until he or she had categorized the object as a snake or a stick. Perhaps the main contribution of Bruner’s early work was highlighting some of the factors that make individuals more likely to interpret information in terms of one category as opposed to another.

According to Bruner, the central factor in determining which category individuals use to interpret information was the relative accessibility of the relevant categories (i.e. the ease with which the individual could retrieve the category). The greater the accessibility of a category the less the stimulus input needed for categorization to occur in terms of that category, the wider the range of input characteristics accepted as belonging in the category,

and the more likely it was that categories that provide a better or equally good fit for the input would be masked. Category accessibility was assumed to be a function of the expectancies of the perceiver and the search requirements imposed by the perceiver’s processing objectives. Thus, a dark tubular shape would be more likely to be categorized as a stick if the perceiver were walking through the snow looking for firewood than if he or she were wading through a jungle stream looking for zoo specimens.

In sum, Bruner proposed that stimulus information was generic until it had been interpreted in terms of a mental category. When individuals have more than one category they could use to interpret information, they use the one whose accessibility had been increased by their motivational state (e.g. hunger) and/or their expectancy (e.g. being in an orchard).

The Classic Social Cognition View: Accessibility X Applicability

Current interest in category accessibility in social psychology can be traced to a study by Higgins, Rholes, & Jones (1977). The initial interest of these investigators was in understanding the way in which the accessibility of trait concepts could influence a person’s interpretation of behaviors as he or she attempted to form an impression of another person. The assumption was that the implications of any given behavior (e.g. skydiving) would depend on the concept (e.g. adventurous versus reckless) used to interpret that behavior. They also assumed, following Bruner, that when behavioral information was interpretable in terms of more than one concept, individuals would use the one that was most accessible. Unlike Bruner, however, Higgins, et al. did not emphasize the role of motivation and expectancy in heightening accessibility. Rather, they emphasized the role of previous activation, which they referred to as priming. They proposed that if a concept had been recently used for almost any processing whatsoever (i.e. if it had been primed), then this concept would be more accessible and thus be more likely to be used to interpret subsequently encountered target information – provided the concepts were applicable to that information. Applicability, in this case, referred to denotative similarity.

To test this passive priming X applicability hypothesis, Higgins, Rholes, & Jones (1977) had participants perform in what, ostensibly, were two unrelated experiments. In the first, participants had to name the color of ink in which various words were written. Then, in the second, participants were presented with a description of a person and were asked to form an impression of this person. This description contained behaviors that were open to several interpretations. For example, the target person was described as being “well aware of his ability to do things well.” This behavior could be interpreted as either self-confident or conceited.

To prime different concepts, Higgins, et al. embedded different words in the color naming task. For some participants, these words were positive and were descriptively relevant to interpreting the target’s behaviors (e.g. self-confident). For other participants, the words were negative yet descriptively relevant to interpreting the target’s behaviors (e.g. conceited). Other participants were exposed to words that were either positive (e.g. neat) or negative (e.g. listless) but that were not relevant to interpreting the target’s behaviors. Consistent with the passive priming X applicability hypothesis, participants rendered more

favorable impressions of the target following activation of the positive compared to the negative concepts, but only when the concepts were denotatively related to the information in the target paragraph.

These results were consistent with Bruner's formulation in the sense that participants used the most accessible, relevant category to interpret target information. The results extended Bruner's formulation, however, by emphasizing the passive possibilities in priming. In the Higgins, Rholes, & Jones (1977) study, participants used the relevant primed concepts even though they were not motivated to do so, had no reason to expect the primed concepts more than their alternatives, and the primed concepts were relevant to the target information only by virtue of their denotative relatedness. It seemed that merely activating a trait concept was sufficient to increase the likelihood that the concept would be used to interpret information to which it was denotatively related.

This view of priming was fleshed out in a number of subsequent studies. Srull & Wyer (1980), for example, demonstrated that primed concepts are used to interpret target information at the time this information is initially encoded. Priming a concept after participants have interpreted the target information has no effect at all on the subsequently formed impression (see also Wyer & Martin, 1986). Higgins, Bargh, & Lombardi (1985) suggested that if a concept is primed frequently enough it can become chronically accessible. When this occurs, the concept is likely to be used in interpreting information even when the concept has not been recently primed by contextual stimuli. Moreover, mental operations performed on information related to a chronically accessible concept may be performed automatically (Bargh & Thein, 1985). Finally, priming effects are not restricted to mental operations or to pencil-and-paper measures. Priming can also affect overt behavior. Priming the concept *aggressive*, for example, can lead individuals to behave more aggressively (Berkowitz, 1993; Carver, Ganellen, Froming, & Chambers, 1983).

Taken together, these early studies painted a very clear and coherent picture of concept priming. The picture was so compelling in fact that a large number of researchers were motivated to use this conceptualization to help make sense of phenomena in a wide variety of areas. For example, researchers applied this passive priming X applicability view of accessibility to stereotypes (Devine, 1989), attitudes (Fazio, Sanbonmatsu, Powell, & Kardes, 1986), goals (Bargh, 1997), relationships (Baldwin, Carrell, & Lopez, 1990), death concerns (Greenberg, Simon, Pyszczynski, & Solomon, 1992), aggression (Berkowitz, 1993), the answering of questions on surveys (Sudman, Bradburn, & Schwarz, 1996), explanations (Wilson, Hodges, & LaFleur, 1995), and anchoring effects (Strack & Mussweiler, 1997).

The advantage of this explosion in research was that the field gained a great deal of knowledge about priming as well as about phenomena that seemed to have priming as one of its underlying components. The flip side of this explosion, of course, was that researchers also became aware of areas of incompleteness. As Smith, Steward, & Buttram (1992, p. 759) noted, "the familiar conceptualization that categorization is a function of Accessibility X Fit now appears to be inadequate." In the remainder of this chapter, we address these inadequacies. More specifically, we present research showing that (a) concept applicability involves more than the denotative fit between the concept and the target; (b) information can be accessible and applicable yet not used to interpret information; (c) increasing the accessibility of different types of knowledge structures (e.g. traits versus exemplars)

produces different effects; and (d) priming can do more than increase the likelihood that a concept will be used to interpret information. We also examine the role of awareness, accuracy motivation, and correction processes in determining the use of primed information.

Non-motivational Qualifications of the Passive Accessibility X Fit View

When distilled to its essence, the passive accessibility X fit view can be seen to consist of five elements: *trait concepts* are used to *interpret* information to the extent that these concepts are *accessible* and *applicable* with applicability being defined in terms of the *denotative similarity* between the primed concept and the target information. In the following sections, we describe qualifications on each of these five elements.

Beyond denotative applicability

As just noted, the initial work on concept accessibility (e.g. Higgins, Rholes, & Jones, 1977) emphasized the denotative fit between the primed knowledge and the target information as a determinant of the likelihood that a primed concept would be used to categorize target information. Subsequent research has suggested, however, that denotative relatedness may not be the only determinant of concept applicability. Stapel & Koomen (1999), for example, had participants form an impression of an ambiguously described target person. As in Higgins, Rholes, & Jones (1977), some participants were primed with concepts that were either denotatively applicable or inapplicable to interpreting the subsequent target information, and that had moderate evaluative connotations (e.g. applicable: assured versus arrogant; inapplicable: unrealistic versus idealistic). With these primes, the results would presumably parallel those of Higgins, Rholes, & Jones (1977). Participants would assimilate their impressions of the target behavior toward the implications of the applicable concepts but not the inapplicable ones.

What would happen, though, if participants were primed with concepts that had no clear denotative implications but that had relatively broad and strong evaluative implications (good versus bad)? Similarly, what would happen if participants were primed with concepts that were denotatively inapplicable to the target information but that had relatively narrow but strong evaluative implications (aggressive versus sweet)? Stapel and Koomen hypothesized that in these cases participants might use the strong evaluative implications to help disambiguate the target information, even though the denotative applicability is low. If so, then participants should assimilate their impression of the target toward the implications of the primed information even though this information is not denotatively relevant to interpreting the target behaviors.

The results of Stapel & Koomen (1999) supported this reasoning. They found that participants' impressions were assimilated toward the implications of the primed concepts not only when these were denotatively applicable, but also when they were generally evaluative with no clear denotative implications and when they were denotatively inapplicable

but possessing strong evaluative implications. It is only when the primed concepts were denotatively inapplicable and both relatively narrow and evaluatively weak that participants' impressions did not assimilate toward the implications of the primed concepts (e.g. Higgins, Rholes, & Jones, 1977). What these results suggest is that denotative applicability is not a necessary ingredient of concept applicability (see also Martin, 1986; Martin, Seta, & Crelia, 1990). In some cases, strong connotation may be sufficient to allow target information to be interpreted in a manner consistent with the implications of primed information, even when this information consists of denotatively inapplicable concepts.

Different types of representations produce different effects

Another feature of the passive priming X applicability view was its emphasis on the use of trait concepts to interpret target information. We know, however, that individuals possess knowledge structures other than trait concepts. These include exemplars, scripts, and procedural knowledge. Subsequent research has begun to explore the effects of priming these knowledge structures. The question is whether effects like those obtained with the priming of trait concepts would also be obtained if one of these other knowledge structures were activated. There is some reason to think not.

Consider, for example, that trait concepts represent diffuse semantic information that can be applicable to a wide range of behaviors. The trait "aggressive," for example, could be used to characterize behaviors as diverse as shoving, verbal abuse, and cutting someone off in traffic. An exemplar, on the other hand, reflects knowledge about a specific person (e.g. Hitler), and this may be less likely to generalize to thoughts about another specific person. An exemplar, however, might make a useful standard of comparison (e.g. Herr, Sherman, & Fazio, 1983). For example, almost anyone would appear less aggressive compared to Hitler. It is possible, therefore, that with the same degree of denotative applicability, priming a trait concept may give rise to assimilation, whereas priming an exemplar may give rise to contrast. These hypotheses were tested, and supported, by Stapel, Koomen, & Van der Pligt (1997). They found that participants primed with trait concepts assimilated their impressions of the target toward the implications of the primed concepts, whereas participants primed with exemplars contrasted their target impressions away from the implications of these exemplars.

If it is true that exemplar priming produces contrast because exemplars make good standards of comparison, then it should be possible to eliminate exemplar-induced contrast by undermining the comparison relevance of the exemplar. This hypothesis was tested by Stapel, Koomen, & Van der Pligt (1997). They primed participants with hostility-related but *non-person* exemplars such as "Shark" and "Tiger" or "Puppy" and "Bunny." These exemplars reflect distinct entities associated with varying degrees of hostility, making them, at least in principle, good candidates for standards of comparison. The exemplars also represent animals, however, and animals are generally not relevant standards of comparison when judging humans. As a result, priming of these non-person exemplars may not lead to contrast of the target person.

These exemplars, however, are associated with different levels of hostility. So, it is possible that when this hostility-related information is made accessible, the result may be

assimilation. That is, participants primed with non-person exemplars may assimilate their impression of an ambiguously described hostile/friendly target person toward the implications of the primed (non-relevant) hostile or friendly exemplars. These were in fact the results obtained by Stapel, Koomen, & Van der Pligt (1997).

In sum, it appears that with the same degree of denotative applicability, priming a trait concept or a non-relevant exemplar is likely to give rise to assimilation, whereas priming a relevant exemplar is likely to give rise to contrast. The former two knowledge structures appear to influence participants' interpretations of the target information, whereas the latter one seems to facilitate a comparison process.

Different effects of primes at encoding versus output

If traits and exemplars play different roles in the impression formation process (i.e. interpretation versus comparison), then priming these different knowledge structures at different times should produce different effects. Recall Srull & Wyer's (1980) finding that priming trait concepts before participants had interpreted the target information resulted in assimilation, whereas priming these same concepts after participants had interpreted the target information had no effect on participants' judgments (see also Wyer & Martin, 1986). What these results suggest is that primed traits are used to disambiguate target information at the time participants first encode that information. When are exemplars used? According to Stapel, Koomen, & Van der Pligt (1997), exemplars may be used as a standard of comparison either at encoding or after individuals have formed an impression and are attempting to translate it into an overt response. What this means is that priming of a trait concept is likely to produce assimilation if the priming occurs before, but not after, participants have encoded the target information, whereas the priming of an exemplar is likely to produce contrast regardless of whether it is primed before or after participants have encoded the information.

To test these ideas, Stapel, Koomen, & Van der Pligt (1997) primed participants with either trait concepts or exemplars, and this occurred either before or after participants had read the target information. In replication of Srull & Wyer (1980), they found assimilation when a trait concept was primed before but not after participants had read the target information. With the priming of exemplars, however, they found contrast regardless of whether the priming occurred before or after participants had read the target information. This pattern of results suggests that trait concepts are generally used in interpreting or disambiguating information as individuals first encode it, whereas exemplars are used as standards of comparison either as individuals are interpreting information or after they have already formed their impression.

The role of procedural knowledge

Another feature common to most early priming research was the tendency to explain priming effects in terms of what might be called structural changes in semantic knowledge. The synapse model (Higgins, et al., 1985), for example, suggested that use of a concept in-

creases its hypothetical charge, which, in turn, made the concept more likely to be used. The bin model (Wyer & Srull, 1989) suggested that use of a concept causes that concept to get placed, metaphorically, on the top of a semantic bin in memory. Because of this placement, the concept would be encountered more quickly than a less recently used concept in any subsequent search of that bin. This, in turn, would make a recently activated concept more likely to be used. As can be seen, despite their differences, both models assumed that priming effects were the result of a change in the status of semantic information. Smith & Branscombe (1987) suggested an alternative. They proposed that at least some category accessibility effects might reflect the operation of procedural knowledge.

Procedural knowledge can be thought of as cognitive structures that represent skills or "how to" knowledge. Such knowledge can be represented hypothetically as production systems or if-then statements. These systems are selected for execution when their conditions (i.e. the "if") match the current contents of working memory or the perceptual environment. The execution of the action part of a production system (i.e. the "then") can result automatically in the performance of cognitive tasks, such as generating inferences. For example, a person may have a production system that specifies something like the following:

IF you observe one *<person>* *<exert power over>* *<another>*,
THEN interpret that behavior as hostile.

The variables (i.e. *person*, *exert power*, *another*) within the production system become instantiated with the values in a given situation. So, if an individual has practiced the production system described above, and if this individual has just observed Donald giving orders to Jamal, then the individual is likely to interpret Donald's action as hostile even though the action might simply reflect Donald's attempt to be efficient. Production systems are assumed to develop out of practice. The basic assumption is that people typically get better at doing things they do frequently.

It is interesting in this context to note that in many priming studies, the priming tasks gave participants repeated practice at interpreting behaviors in terms of a trait. For example, it was not uncommon to prime participants by asking them to construct meaningful sentences out of scrambled words (e.g. "hit he the it"). In performing this task, participants may gain practice in generating trait based interpretations (e.g. hostile) of ambiguous behavior. Could it be that practice of this interpretation procedure (rather than a change in the activation status of concept in semantic memory) was responsible for the subsequent assimilation of the target information? In other words, when presented with a target who engaged in behavior that was relatively ambiguous with regard to its level of hostility, participants may have used the interpretational procedure they had just practiced in the priming task to interpret the target behavior in terms of a hostile concept.

How can we tell if any given priming effect is the result of procedural knowledge or changes in the accessibility of semantic knowledge? According to Smith and Branscombe, the effects of procedural knowledge are more specific and longer lasting than those of semantic priming. To test these hypotheses, Smith & Branscombe (1987, Experiment 2) had participants perform a task that either allowed them to practice the procedure of interpreting behavior in terms of a concept or that increased the accessibility of semantic

knowledge. In the first case, participants were asked to construct meaningful sentences from a scrambled list of words. In the second case, participants were asked to judge pairs of words (e.g. hostile-crude) in terms of whether or not the words had the same meaning. Then, either immediately following these tasks or three minutes later, participants were presented with a behavioral description that was ambiguous with regard to the trait (or procedure) that had been primed in the earlier tasks.

Consistent with the hypothesis that procedurally mediated priming effects last longer than semantically mediated priming effects, Smith and Branscombe found that when the priming task consisted of unscrambling sentences, participants' impressions assimilated toward the implications of the primed concepts in both the short and the long delay. When the priming task consisted of judging the meaning of words, participants' impressions assimilated toward the implications of the primed concepts only in the short delay condition.

In a second study, Smith and Branscombe found some evidence consistent with their second proposed distinction between procedurally mediated and semantically mediated priming effects. They found that the effects of procedural priming are more specific than those of semantic trait priming. Taken together, the results of these two studies suggest that the activation of procedural knowledge can account for at least some priming effects, and that theoretical conceptualizations of priming that focus only on the activation of semantic concepts are incomplete.

Priming does more than facilitate categorization

We have seen that information other than traits (e.g. exemplars, procedures) can be primed and that effects other than interpretation (e.g. contrast) can occur. Such results suggest that priming does more than increase the probability that individuals will use primed trait concepts to interpret behaviors. It appears, instead, that priming increases the likelihood that the primed knowledge will be used in whatever processing occurs at the time, whether this be interpretational or otherwise. This more general view of priming may be best illustrated in research showing that priming can influence quantitative judgments that do not even involve interpretation. Strack & Mussweiler (1997), for example, used priming to explain the anchoring and adjustment effect.

In one demonstration of the anchoring and adjustment effect, Tversky & Kahneman (1974) asked participants whether the percentage of African nations in the United Nations was higher or lower than 80 percent (or 20 percent). Then, they asked these same participants to estimate the actual percentage of African nations in the United Nations. Tversky and Kahneman found that participants given the high standard (e.g. 80 percent) in the comparison task provided higher estimates on the absolute judgment task than did participants given the low standard (e.g. 20 percent).

Explanations of this effect have focused on the numerical value provided as a standard in the initial comparison task. More specifically, it has been suggested (e.g. Jacowitz & Kahneman, 1995) that participants start their estimation of the absolute value at the value presented in the comparison task. Then, they adjust upward or downward (depending on the condition) until they encounter the outer limits of their acceptable range of responses.

This results in higher estimates when participants are adjusting downward than when they are adjusting upward.

Strack & Mussweiler (1997) explored the possibility that performing the initial comparison task might do more than leave a numerical standard salient. It might also prime a more general mental representation relevant to rendering the subsequent absolute judgment. Specifically, Strack and Mussweiler proposed that in performing the comparison task, participants might create a representation of the target as possessing the standard. For example, if participants were asked whether the Mississippi River is longer or shorter than 3,000 miles, they might imagine the north-south extension of the United States and use their geographic knowledge to compute the answer. If they do this using a positive test strategy (Klayman & Ha, 1987), then they are likely to bring to mind information consistent with the target possessing the value supplied in the standard. This standard-consistent information may then be consulted when participants make the subsequent absolute judgment, and it is this information that leads to the anchoring and adjustment effect.

Note that this standard-consistent information would not be brought to mind if the standard provided by the comparison task were clearly implausible (e.g. "Is the Mississippi River longer or shorter than 1 mile?"). In this case, the standard would be so clearly wrong that participants could generate an answer to the comparison question without bringing a great deal of related information to mind. If this theoretical analysis is correct, then participants should take longer to make the initial comparison judgment when a plausible rather than implausible standard is provided.

The effects of standard plausibility should be different, however, on the time it takes participants to render the absolute judgment. Because participants presented with a plausible standard would already have brought related information to mind, they should perform the absolute judgment relatively quickly. Participants presented with an implausible standard, on the other hand, will not have brought such information to mind, so they must do so when they make the absolute judgment. This could take time. This reasoning implies a crossover interaction in response times. Relative to participants who have been primed with an implausible standard, those primed with a plausible standard will take more time to make the initial comparison judgment but less time to make the subsequent absolute judgment. This crossover pattern was in fact obtained by Strack and Mussweiler.

If, as these results suggest, informational priming plays a role in (at least some) anchoring and adjustment effects, then these effects might not be seen if the primed information were inapplicable to the absolute judgment. This is because primed information is applied only to the extent it is applicable (Higgins, Rholes, & Jones, 1977; Stapel & Koomen, 1999). To test this idea, Strack and Mussweiler asked participants to render a comparison judgment (e.g. "Is the Brandenburg Gate taller or shorter than 50 (or 150) meters?") followed by an absolute judgment. The absolute judgment was either related to the information primed in the comparison task (e.g. "How tall is the Brandenburg Gate?") or it was not (e.g. "How wide is the Brandenburg Gate?"). Consistent with the priming applicability hypothesis, an anchoring and adjustment effect was observed only when the information primed in the comparison task was applicable to the absolute judgment (e.g. both judgments involved height).

Taken together, these studies suggest that at least some anchoring and adjustment effects are not due simply to participants starting from a numerical standard and then

adjusting upward or downward to an insufficient degree. Rather, priming of more general, judgment relevant information seems to be involved. This primed information is not a trait concept, however, and it is not used to interpret target information. It is a mental representation of the target possessing the value primed in the standard, and it can influence subsequent quantitative judgments to the extent that it is applicable to those judgments.

Content or Phenomenology?

The studies we have discussed so far were concerned primarily with the nature of the content that had been made accessible. These studies examined the effects of priming concepts, exemplars, or procedures that varied in their applicability to the target information. In each case, the general assumption was that target judgments are affected because of the implications of the content that has been brought to mind. Note, however, that priming not only brings information to mind, it also increases the ease with which this information comes to mind. Could the subjective experience of ease of retrieval also be informative? This possibility was first raised by Tversky & Kahneman (1974) in their discussion of the availability heuristic. They suggested that individuals sometimes assess the frequency of an occurrence (e.g. words beginning with the letter k) by assessing the ease with which they can retrieve instances of that occurrence.

Schwarz, Bless, Strack, & Klumpp (1991) examined the effects of priming in a condition in which the implications of the primed content were at odds with the implications of subjective ease of retrieval. Specifically, they had participants recall either 6 or 12 instances of either assertive or submissive behaviors they had performed. Then, they had participants rate themselves in terms of how assertive-submissive they were. Participants who recalled 12 instances of the target behavior would have twice as much evidence that they possessed the target trait than participants who only recalled 6 instances. So, we might expect participants who recalled 12 instances to rate themselves higher in the direction of the recalled instances than participants who recalled only 6. This would reflect the impact of information accessibility.

On the other hand, it is more difficult to recall 12 instances than 6 instances. So, if participants used ease of retrieval to estimate frequency, then we might expect those who recalled 12 instances to rate themselves as possessing less of the trait than participants who recalled only 6 instances. If participants really were assertive (or submissive), then why would it be so hard for them to recall 12 instances in which they displayed this trait? Or so their thinking would go.

The results suggested that both accessibility and ease of retrieval were informative. When participants recalled only 6 instances, those who recalled the assertive behaviors rated themselves as more assertive than those who recalled the submissive behaviors. When participants recalled 12 instances, however, those who recalled the assertive behaviors rated themselves as less assertive than those who recalled the submissive behaviors. In short, individuals can gain information not only from what comes to mind but also from *how* that information comes to mind.

Motivated Limits on the Passive Priming X Fit View

What we have seen so far is that the familiar conceptualization of accessible, applicable trait concepts being used to interpret information can be qualified in each of its major components. Different knowledge structures activated at different times can have different effects. Note, however, that each of the qualifications we discussed were non-motivational. That is, they were based on the nature of the knowledge structure primed and/or the time at which the structures were primed. There is evidence, however, that motivational factors can also determine the nature of priming effects. We can take a trait concept with a given level of accessibility and applicability, for example, and influence the extent to which participants are likely to use that concept to interpret information by manipulating variables such as accuracy motivation (Thompson, Roman, Moskowitz, Chaiken, & Bargh, 1994) and cognitive effort (Martin, Seta, & Crelia, 1990). One implication of such findings is that a distinction needs to be maintained between a concept's accessibility and its use in interpreting information (Martin, 1986).

To the extent that a concept's accessibility and its use are distinct, it should be possible to manipulate the two orthogonally. This was accomplished by Stapel & Koomen (under review). They hypothesized that priming a concept would lead to different judgments among participants who approached the judgment task with an interpretational mindset compared to participants who approached the task with a comparison mindset. To induce an interpretational set, Stapel and Koomen presented participants with a list of different behaviors and asked the participants to describe each behavior with a single word. To induce a comparison mindset, participants were asked to compare the persons performing the behaviors to different standards (e.g. the average woman, the average student). Following this task, participants were primed with either a positive or a negative concept and asked to form an impression of a target whose behavior was relatively ambiguous with regard to the primed traits.

Consistent with the idea that priming of the same trait can have different effects on judgments of the same behavior depending on one's mindset, participants who had completed the interpretation task assimilated their impressions toward the implications of the primed concepts, whereas participants who had completed the comparison task contrasted their impressions with the implications of the primed concepts. Thus, motivational variables can influence priming effects even when the structural properties of the situation (e.g. the applicability and accessibility of the primed concepts) are held constant.

Epistemic motivation

Within social psychology, motivation has often been identified with the irrational, such as dissonance or ego-defensive biases, and it has typically been contrasted with cold, analytical, logical thinking such as that typified in attribution theory. From a general information processing perspective, however, there can also be motivations toward accuracy and analytic thinking. Moreover, these kinds of motivations can moderate the effects of concept priming.

Ford & Kruglanski (1995) and Thompson, et al. (1994), for example, found that participants with no particular incentive to be accurate while forming an impression assimilated their impression of an ambiguous target person toward the implications of the primed traits. The impressions of participants who were highly motivated to form an accurate impression, on the other hand, showed no influence of the primed concepts. Apparently, the motivation to be accurate can, at least in some cases, inoculate participants from the assimilative effects of concept priming.

In a conceptually related study, Sedikides (1990) found that the more specific motivation of communicating a particular impression can also attenuate the assimilative effects of concept priming. Participants who thought they were to communicate their impression of an ambiguously described target person to an audience that liked the target rendered more favorable impressions of the target than participants who thought they had to convey their impression to an audience that disliked the target. More importantly, for present purposes, these differences in judgment occurred even when participants were primed with concepts that were opposite in valence to the presumed attitudes of the audience. In other words, participants communicating an impression to an unfavorable audience formed a negative impression despite being primed with a positive, applicable concept. Taken together, these studies suggest that certain processing objectives (e.g. be accurate, communicate a specific impression) can override the effects of passive contextual priming.

Not all priming effects are susceptible to moderation by accuracy motivations, however. Stapel, Koomen, & Zeelenberg (1998) found that accuracy motives were less likely to overcome contrastive influences than assimilative influences. These investigators induced assimilation or contrast by priming a trait concept or an exemplar, respectively (cf. Stapel, Koomen, & Van der Pligt, 1997). Following this, they had participants form an impression of an ambiguously described target person. To manipulate accuracy motivation, half of the participants were told that the tasks they were performing were merely part of a pilot study and might possibly be used in some future study. The remaining participants were told that they should try to be as accurate as possible in their judgments. In replication of earlier work, trait priming led to assimilation among low accuracy participants, but had no effect among the high accuracy participants. Unlike earlier studies, however, exemplar priming led to contrast in both the low accuracy and the high accuracy participants. Taken together, the studies discussed in this section suggest that different motivations have different consequences for different types of knowledge accessibility effects.

Awareness of a bias

Another type of motivation that has been studied is more specific, namely, removing perceived bias from the target judgment. Note that the initial priming research (e.g. Higgins, Rholes, & Jones, 1977) emphasized the passive nature of priming. This was true, in part, because the initial studies used disguised priming tasks in the so-called "two-experiment" paradigm. These steps were taken to rule out the possibility that the results were due to demand characteristics. They were successful in doing this, but they may also have ruled out the possibility that participants would engage in other theoretically important psychological processes. As Martin (1986, p. 494) noted:

When a concept is primed very subtly, individuals may not even be aware that it has been activated in them (Bargh & Pietromonaco, 1982). Consequently, when this concept comes to mind in the subsequent impression formation task they have no reason to believe that it is anything other than their own spontaneous reaction to the target. This means that they have no reason not to use the primed concept in interpreting the target information, provided that it is consistent with the implications of that information.

The same may not hold true, however, when a concept is primed more blatantly. Under these conditions, individuals may associate the activation of the concept with their exposure to the priming stimuli rather than with the target stimulus. As a result, they may actually avoid using the primed concept to interpret the target information, as its use would appear to bias their independent evaluation of the target.

In short, a quite different set of processes may come into play when concepts are primed more blatantly, and these processes may produce quite different effects than those found in the initial priming studies. Evidence for this possibility has been obtained in a variety of studies. Strack, Schwarz, Bless, Kübler, & Wänke (1993), for example, had participants perform what they thought were a series of cognitive and perceptual tasks. In one of these tasks, participants heard a series of tones paired with words. Participants were asked to classify the tones as high or low and to write down the words. For half of the participants, the words were positive (e.g. friendship), whereas for the other half, the words were negative (e.g. dishonest). Following the tone–word task, participants were asked to form an impression of a person whose actions (e.g. stole exam questions for a desperate friend) were interpretable in terms of either the positive or the negative primed concepts.

To manipulate the blatancy of the priming stimuli, Strack, et al. (1993) had some participants perform the tone–word task and then form their impression. Other participants were asked to answer some questions about the tone–word task (e.g. how well they were able to discriminate the tones) prior to forming their impressions. The point of this questioning was to remind participants of the priming stimuli (i.e. the positive or negative words used in the tone–word task).

In replication of earlier priming studies (e.g. Higgins, Rholes, & Jones, 1997), Strack, et al. (1993) found that impressions of the target assimilated toward the primed concepts when participants were not reminded of the priming task. When participants were reminded, however, their impressions were contrasted with the implications of the primed concepts. What these results suggest is that a concept may be highly accessible and applicable to interpreting information, but if participants are aware that this concept has been primed by a non-target event (i.e. the priming stimuli), then they may not use this concept in interpreting the target information (see also Lombardi, Higgins, & Bargh, 1987). The next question, of course, is why. Why do participants avoid the use of blatantly primed concepts?

Several explanations have been offered (Lombardi, Higgins, & Bargh, 1987; Martin, 1986; Strack & Hannover, 1996; Wegener & Petty, 1995; Wilson & Brekke, 1994). Despite their differences, these explanations generally agree that, in some way or another, participants sense a threat to the genuineness of their evaluation of the target and (in some form or another) they take steps to remove the perceived bias from their evaluation of the target. So, now we can ask, what is the nature of the bias that is sensed? How do people correct for this bias? One possible answer to both questions entails naive theories, that is, beliefs about the effects of a context on a target.

Theory Based Correction

According to Wilson & Brekke (1994), individuals generally have weak introspective abilities (see also Nisbett & Wilson, 1977). One implication of this weakness is that individuals may generally fail to appreciate the influence of contextual factors on their judgments. In other words, individuals may not be able to discriminate reliably between biased and unbiased judgments merely by turning inward. So, if they are to remove the bias from their judgments, then they must look elsewhere for guidance.

One source of guidance may be naive theories. Consider, for example, a person who is asked to rate a moderately attractive face after having just rated some extremely attractive faces. In this context, the moderately attractive face may appear to the person to be unattractive (i.e. a contrast effect). Although this is a biased judgment, the person may not realize it. The assessment of the face as unattractive may feel like the person's genuine assessment of the target. Suppose, however, that the person retrieved a theory that suggested that ratings of moderately attractive faces could be lowered by previously rating more attractive faces. Armed with this knowledge, the person would be in a position to correct for the biasing influence of the context. In this way, individuals' naive theories could potentially alert them to biases that they might miss if they relied solely on introspection.

It appears, further, that naive theories may help people to remove the perceived bias from their judgments. Specifically, it has been suggested (Strack, 1992; Strack & Hannover, 1996; Wegener & Petty, 1995; Wilson & Brekke, 1994) that when people believe that their judgments are being biased, they consult their naive theories to determine the extent and direction of the bias. Then, they adjust their target ratings in a direction that is opposite to the theorized bias and to an extent that is commensurate with the theorized amount of bias. As Strack (1992, p. 269) put it, "People can apply norms, rules, or theories to adjust their response for the effect of the pernicious influence. . . . It is important, however, that judges have such rules at the ready; otherwise, they would not know how to alter their responses."

Evidence suggestive of a role for naive theories in the correction of contextual bias has been obtained in a series of studies by Wegener and Petty (for a review, see Wegener & Petty, 1997). They began by providing participants with a series of context–target configurations and asking participants to indicate what effect the context might have on ratings of the target. In this way, Wegener and Petty were able to find sets of stimuli for which participants held theories of either assimilation or contrast. For example, most participants believed that their ratings of a product would be biased toward desirability if the product were endorsed by attractive as compared to unattractive women (i.e. an assimilation effect). Most also believed that ratings of moderately attractive women would be biased away from ratings of extremely attractive or extremely unattractive women (i.e. a contrast effect).

After establishing that there were sets of stimuli for which participants held theories of either assimilation or contrast, Wegener & Petty (1995) had participants actually make their ratings of these stimuli. Half of the participants were asked without further elaboration to rate the context and target items, whereas half were given an explicit warning not to

let their judgments of the context influence their judgments of the target. This warning informed participants of a possible bias, but it did not specify the direction or magnitude of that bias. It was assumed that this information would be gleaned by participants from their naive theories. The results were consistent with this hypothesis.

When participants simply rated the context and target stimuli, their target judgments reflected assimilation when participants rated stimuli they had earlier theorized would lead to assimilation, but reflected contrast when they rated stimuli they had earlier theorized would lead to contrast. When participants had been instructed to remove the contextual bias, however, their judgments showed the opposite pattern. There was assimilation when participants rated stimuli they had earlier theorized would lead to contrast, but contrast when they had rated stimuli they had earlier theorized would lead to assimilation. This pattern is consistent with the hypothesis that when individuals are alerted to a potential bias in their judgments, they consult their naive theories in order to understand the nature of the contextual influence, and then they adjust their judgments in a direction opposite to the theorized influence.

Target Based Correction

According to the theory based models of correction, individuals cannot correct for a contextual influence without having some sort of understanding of the nature of that influence. Although this assumption is plausible, it is not entirely clear that the knowledge individuals need to correct their judgments has to come from a theory that specifies the context–target relation. As Stapel, Martin, & Schwarz (1998) noted, the blatant warning used in the Wegener and Petty studies may have allowed participants to by-pass their use of theories. The warning (i.e. “Please don’t let your ratings of the target be influenced by your rating of the context”) essentially told participants that the contextual stimuli were likely to be biasing their target judgments. As a result, participants did not have to consult their theories to determine if the context were biasing their judgments. They already had reason to believe it was. So, a blatant warning may allow participants to by-pass the first step of theory-guided correction (i.e. use theory to detect bias).

Of course, even if participants do not use their theories to detect the contextual bias, then they may still consult their theories when correcting for the bias. After all, participants still need to know the direction and extent of the influence if they are to correct for it. Stapel, Martin, & Schwarz (1998) hypothesized, however, that participants may even be able to correct without recourse to a theory that specifies the relation between the context and the target. If participants experience an inclination to evaluate the target favorably and are told that there may be a bias in their judgments, then participants may infer that an unbiased judgment would be one that was less favorable. Conversely, if they experience an inclination to evaluate the target unfavorably and are told that there may be a bias in their judgments, then participants may infer that an unbiased judgment would be one that was more favorable. In this way, participants’ evaluative inclinations in the context of a blatant warning may allow participants to adjust their judgments without consulting theories to assess bias or guide correction.

Stapel, Martin, & Schwarz (1998) addressed this issue by examining the nature of correction induced by blatant and subtle warnings. All participants were asked to rate some target stimuli in a context that typically produced contrast (e.g. they rated the desirability of the weather in Midwestern US cities after having rated the desirability of the weather in vacation spots). One group of participants was asked merely to rate the context and target stimuli. Two other groups rated these same stimuli, but received a warning between their ratings of the context and the targets. For some participants, this warning was the blatant one used by Wegener and Petty. Specifically, these participants were instructed to “Make sure that your perceptions of the weather in the vacation spots above do not influence your ratings of the following places.” The remaining participants received a conditional warning. They read “When you feel there is something that may have an unwanted influence on your ratings, please try to adjust for that influence.” The first warning implies that there is a bias; the second allows participants to determine on their own whether or not there is a bias.

The next step was to manipulate the amount of bias participants perceived to be coming from the context. This was accomplished by having half of the participants rate the context and target stimuli on the same dimension, but having half rate the context on one dimension and the targets on another dimension (e.g. job satisfaction versus desirability of weather). The potential biasing influence of the context should be more obvious when the stimuli are considered on the same dimension than when they are considered on different dimensions (cf. Brown, 1953). So, if participants correct only when they perceive a bias, then they are likely to correct only when they rate the context and target on the same dimension.

The results indicated a clear difference between the effects of the subtle and blatant warnings. When participants received no warning, their judgments reflected contrast regardless of whether they rated the targets on the same dimension as the contextual stimuli or on different dimensions. This suggests that the uncorrected effect of this context was contrast. What this also means is that if the warned participants correct for the contextual influence, then their judgments will shift toward assimilation. When participants were given the conditional warning, such a shift toward assimilation was seen, but only when the perceived influence of the context was obvious.

Judgments of participants who had been blatantly warned, on the other hand, reflected a correction toward assimilation regardless of the dimension on which the contextual stimuli had been judged. In other words, the blatantly warned participants adjusted their responses regardless of whether the level of bias coming from the contextual stimuli was salient or non-salient. More importantly, the blatantly warned participants adjusted their target evaluations even in a condition in which the conditionally warned participants did not perceive any bias. If the conditionally warned participants did not detect a bias coming from this context, then what were the blatantly warned participants correcting for?

These data raise the possibility that correction may take place without reference to a theory that specifies the context–target influence. Participants may simply consider their target judgment in the context of a blatant warning. If they experience an inclination to render a favorable judgment, then they might make their judgments less positive. If they experience an inclination to render an unfavorable judgment, then they might make their judgments more positive.

Summary and Conclusions

The early priming research depicted the processes involved in priming as rather simple. It was generally assumed that individuals used whatever concepts were accessible and applicable with applicability being defined in terms of denotative similarity. Subsequent research, however, has qualified this view. In fact, it appears that this view of priming holds only under the conditions used in the initial studies. When changes, even relatively small ones, are made to the procedures used in the early studies, quite different results are obtained. The recent studies highlight the sophisticated and conditional nature of priming effects. These effects depend on the type of knowledge structure activated, the timing of the activation, and the motivation level and mental set of the participants.

Although social judgments that have been influenced by priming are not always accurate in an objective sense, this should not take away from the function of priming, which is to fine-tune individuals' processing to the specific judgment task at hand. It is in the nature of judgment that individuals bring aspects of themselves (memories, processing objectives) to bear and these aspects influence the individuals' judgments in predictable ways. We should continue to explore the regularities governing the construction of social judgments and we should give the social perceiver his or her due as sophisticated processors of social information.

References

- Baldwin, M. W., Carrell, S. E., & Lopez, D. F. (1990). Priming relationship schemas: My advisor and the pope are watching me from the back of my mind. *Journal of Experimental Social Psychology*, 26, 435-454.
- Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer, Jr. (Ed.), *Advances in social cognition* Vol. 10 (pp. 1-61). Mahwah, NJ: Lawrence Erlbaum Associates.
- Bargh, J. A., & Pietromonaco, P. (1982). Automatic information processing and social perception: The influence of trait information presented outside of conscious awareness on impression formation. *Journal of Personality and Social Psychology*, 43, 437-449.
- Bargh, J. A., & Thein, R. D. (1985). Individual construct accessibility, person memory, and the recall-judgment link: The case of information overload. *Journal of Personality and Social Psychology*, 49, 1129-1146.
- Berkowitz, L. (1993). Towards a general theory of anger and emotional aggression: Implications of the cognitive-neoassociationistic perspective for the analysis of anger and other emotions. In R. S. Wyer, Jr., & T. K. Srull (Eds.), *Advances in social cognition* Vol. 6 (pp. 1-46). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brown, D. R. (1953). Stimulus-similarity and the anchoring of subjective scales. *American Journal of Psychology*, 66, 199-214.
- Bruner, J. S. (1957). On perceptual readiness. *Psychological Review*.
- Bruner, J. S. (1992). Another look at New Look 1. *American Psychologist*, 47, 780-783.
- Carver, C. S., Ganellen, R. J., Froming, W. J., & Chambers, W. (1983). Modelling: An analysis in terms of category accessibility. *Journal of Experimental Social Psychology*, 16, 779-804.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5-18.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229-238.
- Ford, T. E., & Kruglanski, A. W. (1995). Effects of epistemic motivations on the use of accessible constructs in social judgment. *Personality and Social Psychology Bulletin*, 21, 950-962.
- Greenberg, J., Simon, L., Pyszczynski, T., & Solomon, S. (1992). Terror management and tolerance: Does mortality salience always intensify negative reactions to others who threaten one's worldview? *Journal of Personality and Social Psychology*, 63, 212-220.
- Hastorf, A. H., & Cantril, H. (1954). They saw a game: A case study. *Journal of Abnormal and Social Psychology*, 49, 129-14.
- Herr, P. M., Sherman, S. J., & Fazio, R. H. (1983). On the consequences of priming: Assimilation and contrast effects. *Journal of Experimental Social Psychology*, 19, 323-340.
- Higgins, E. T., Bargh, J. A., & Lombardi, W. (1985). Nature of priming effects on categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 59-69.
- Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, 13, 141-154.
- Jacowitz, K. E., & Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychology Bulletin*, 21, 1161-1166.
- Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypotheses testing. *Psychological Review*, 94, 211-228.
- Lombardi, W. J., Higgins, E. T., & Bargh, J. A. (1987). The role of consciousness in priming effects on categorization: Assimilation versus contrast as a function of awareness of the priming task. *Personality and Social Psychology Bulletin*, 13, 411-429.
- Martin, L. L. (1986). Set/reset: The use and disuse of concepts in impression formation. *Journal of Personality and Social Psychology*, 51, 493-504.
- Martin, L. L., & Achee, J. W. (1992). Beyond accessibility: The role of processing objectives in judgment. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 195-216). Hillsdale, NJ: Erlbaum.
- Martin, L. L., & Tesser, A. (1992). *The construction of social judgment*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Martin, L. L., Seta, J. J., & Crelia, R. (1990). Assimilation and contrast as a function of people's willingness and ability to expend effort in forming an impression. *Journal of Personality and Social Psychology*, 59, 27-37.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84, 231-259.
- Schwarz, N., Bless, H., Strack, F., & Klumpp, G. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61, 195-202.
- Sedikides, C. (1990). Effects of fortuitously activated constructs versus activated communication goals on person impressions. *Journal of Personality and Social Psychology*, 58, 397-408.
- Smith, E. R., & Branscombe, N. R. (1987). Procedurally mediated social inferences: The case of category accessibility effects. *Journal of Experimental Social Psychology*, 23, 361-382.
- Smith, E. R., Steward, T. L., & Buttram, R. T. (1992). Inferring a trait from a behavior has long-term, highly specific effects. *Journal of Personality and Social Psychology*, 62, 753-759.
- Srull, T. K., & Wyer, R. S. (1979). The role of category accessibility in the interpretation of information about persons: Some determinants and implications. *Journal of Personality and Social Psychology*, 37, 1660-1672.
- Srull, T. K., & Wyer, R. S. (1980). Category accessibility and social perception: Some implications for the study of person memory and interpersonal judgments. *Journal of Personality and Social Psychology*, 38, 842-856.
- Stapel, D. A., & Koomen, W. (1997). Using primed exemplars during impression formation: Interpretation or comparison? *European Journal of Social Psychology*, 27, 357-367.
- Stapel, D. A., & Koomen, W. (in press). How far do we go beyond the information given? The impact of knowledge activation on interpretation and inferences. *Journal of Personality and Social Psychology*.
- Stapel, D. A., & Koomen, W. (under review). The impact of interpretation versus comparison goals on knowledge accessibility effects. Manuscript under review.
- Stapel, D. A., Koomen, W., & Van der Pligt, J. (1997). Categories of category accessibility: The

- impact of trait concept versus exemplar priming on person judgments. *Journal of Experimental Social Psychology*, 33, 47-76.
- Stapel, D. A., Koomen, W., & Zeelenberg, M.. (1998). The impact of accuracy motivation in interpretation, comparison, and correction processes: Accuracy X knowledge accessibility effects. *Journal of Personality and Social Psychology*, 74, 878-898.
- Stapel, D. A., Martin, L. L., & Schwarz, N. (1998). The smell of bias: What instigates correction processes in social judgments? *Personality and Social Psychology Bulletin*, 24, 797-806.
- Strack, F. (1992). The different routes to social judgments: Experiential vs. informational strategies. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgment* (pp. 249-275). Hillsdale: Erlbaum.
- Strack, F., & Hannover, B. (1996). Awareness of influence as a precondition for implementing correctional goals. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 579-596). New York: Guilford Press.
- Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73, 437-446.
- Strack, F., Schwarz, N., Bless, H., Kübler, A., & Wänke, M. (1993). Awareness of the influence as a determinant of assimilation vs. contrast. *European Journal of Experimental Social Psychology*, 23, 53-62.
- Sudman, S., Bradburn, N. M., & Schwarz, N. (1996). *Thinking about answers: The application of cognitive processes to survey methodology*. San Francisco: Jossey-Bass Publishers.
- Thompson, E. P., Roman, R. J., Moskowitz, G. B., Chaiken, S., & Bargh, J. A. (1994). Accuracy motivation attenuates covert priming: The systematic reprocessing of social information. *Journal of Personality and Social Psychology*, 66, 474-489.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.
- Wegener, D. T., & Petty, R. E. (1995). Flexible correction processes in social judgment: The role of naive theories in corrections for perceived bias. *Journal of Personality and Social Psychology*, 68, 36-51.
- Wegener, D. T., & Petty, R. E. (1997). The flexible correction model: The role of naive theories of bias in bias correction. In M. P. Zanna (Ed.), *Advances in experimental social psychology* Vol. 29 (pp. 141-208). San Diego: Academic Press.
- Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin*, 116, 117-142.
- Wilson, T. D., Hodges, S. D., & LaFleur, S. J. (1995). Effects of introspecting about reasons: Inferring attitudes from accessible thoughts. *Journal of Personality and Social Psychology*, 69, 16-28.
- Wyer, R. S., Jr., & Martin, L. L. (1986). Person memory: The role of traits, group stereotypes, and specific behaviors in the cognitive representation of persons. *Journal of Personality and Social Psychology*, 50, 661-675.