

Sensitivity and Flexibility

Exploring the Knowledge Function of Automatic Attitudes

MELISSA J. FERGUSON
JOHN A. BARGH

The concept of attitudes is historically one of the most widely studied topics in social psychology (e.g., see Allport, 1935; Eagly & Chaiken, 1993; McGuire, 1986). An attitude has generally been consistently defined throughout the past 7 decades as “a psychological tendency that is expressed by evaluating a particular entity with some degree of favor or disfavor” (Eagly & Chaiken, 1993, p. 1; see also Allport, 1935; Smith, Bruner, & White, 1956; Thompson, Zanna, & Griffin, 1995; Thurstone, 1928). The research activity surrounding this construct has addressed a wide spectrum of topics, such as underlying structure (e.g., Eagly & Chaiken, 1993; Katz & Stotland, 1959; Rosenberg & Hovland, 1960), functionality (e.g., Katz, 1960; Katz & Stotland, 1959; Pratkanis, Breckler, & Greenwald, 1989), the ability of attitudes to predict attitude-relevant behavior (e.g., Elms, 1975; Fishbein & Ajzen, 1974, 1975; Kelman, 1974; Schuman & Johnson, 1976; Wicker, 1969), and methods of attitude measurement (e.g., see Himmelfarb, 1993, for a review).

Though many of these issues are currently being examined, much of the zeitgeist concerns the degree to which the particular measurement of attitudes has consequences for the operationalization and utility of attitudes. More specifically, recent research and theory focus on whether and how attitudes that are activated *outside of awareness* differ

from attitudes that are *strategically recalled* (e.g., Ajzen & Sexton, 1999; Banaji, 2001; Cunningham, Preacher, & Banaji, 2001; Devine, 1989; Devine, Kawakami, Johnson, Johnson, & Howard, 1997; Fazio, Jackson, Dunton, & Williams, 1995; Klauer, 1998; McConnell & Leibold, 2001; Nesdale & Durkin, 1998; Rudman, Greenwald, Mellott, & Schwartz, 1999; Wilson, Lindsey, & Schooler, 2000; Wittenbrink, Judd, & Park, 1997). For example, to what degree can researchers apply theoretical perspectives on the nature of strategically recalled attitudes to automatically activated attitudes?

In this chapter we extend a traditional analysis of strategic attitudes to automatic attitudes. Researchers and theorists have delineated the ways in which strategic attitudes are functional because of their delivery of affective information concerning the corresponding objects (e.g., Katz, 1960; Pratkanis et al., 1989; Smith, Bruner, & White, 1956). We argue that automatic attitudes can be similarly conceptualized as functional in terms of both their sensitivity to contextual differences in the meaningfulness of objects, and their flexibility in responding to novel objects. We first describe strategic and implicit measurement in general and next consider the particular ways in which both strategic and automatic attitudes can be considered functional. We then turn to recent research concerning the sensitivity and flexibility of automatic attitudes.

STRATEGIC RECALL: EXPLICIT ATTITUDE MEASUREMENT

Attitudes have traditionally been measured by simply asking respondents to report their preferences, feelings, thoughts, and behaviors regarding an issue or object. Although such measurement captures the intentional stance of people (i.e., it measures what people consciously and intentionally report), it generates data that are otherwise fraught with interpretational difficulties. Because people can strategically modify their answers, their "attitudes" can represent factors other than the underlying construct of interest (e.g., Devine, 1989, 1995; Fazio et al., 1995; Greenwald, McGhee, & Schwartz, 1998; Jones & Sigall, 1971; Schuman & Kalton, 1985; Warner, 1965). For example, people may, on some occasions, present themselves as holding socially desirable, egalitarian values even though they demonstrate (socially undesirable) racial prejudices on other, less public occasions (e.g., Sigall & Page, 1971). People may report an attitude according to what they believe the experimenter expects (e.g., Orne, 1962; Rosenthal & Jacobson, 1966) or what the social norms prescribe (e.g., Gaes, Kalle, & Tedeschi, 1978; Ostrom, 1973).

The possibility that respondents can modify their attitudes on explicit measures led researchers to develop more covert explicit attitude measures, wherein the attitude object of interest is not obvious. For example, researchers interested in assessing the extent to which people are racially prejudiced constructed subtle measures of racism that ostensibly tap political conservatism (e.g., McConahay, 1986). This masking of the true construct of interest is intended to minimize the degree to which participants react to social norms and demand effects. If participants believe that the questions pertain to policy preferences, they may not hide or modify their true racially prejudiced attitudes.

AUTOMATIC ATTITUDES: IMPLICIT ATTITUDE MEASUREMENT

Even explicit measures that are subtle, however, do not preclude the possibility that respondents strategically misrepresent their feelings or beliefs in the interest of some attitude-irrelevant factor. The possibility that people alter their explicitly reported attitudes means that such measures are perhaps not capturing "real" or unadulterated, attitudes. This skepticism has encouraged researchers to adopt implicit measures of attitudes (e.g., Bargh, Chaiken, Gøvender, & Pratto, 1992; Bargh, Chaiken, Raymond, & Hymes, 1996; Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Greenwald et al., 1989; Greenwald et al., 1998; Wittenbrink, Judd, & Park, 1997). Implicit measures can assess the phenomenon of interest without participants being aware that their reactions to objects are being recorded. Because they are unaware of the intent of the measure, they are theoretically unable to react in a socially desirable or demand-laden fashion (though see Glaser & Banaji, 1999).

Implicit measures of attitudes are currently generating numerous research questions (e.g., Banaji, 2001; Fazio et al., 1986; Greenwald et al., 1998). One of the most frequently used implicit attitude measures (e.g., Bargh et al., 1992; Bargh et al., 1996; Fazio et al., 1986; Fazio et al., 1995) entails a sequential priming paradigm (Neely, 1976, 1977) in which prime words that represent attitude objects (e.g., "puppy," "crime") are paired with positively or negatively valenced adjectives (e.g., "generous," "awful"). Within a typical trial, a prime word is presented for a fraction of a second and is then followed by an adjective, to which participants must respond in some way (e.g., evaluation task, lexical decision task). The response times to the positive and negative adjectives are compared as a function of the valence of the preceding primes. This paradigm demonstrates the phenomenon of evaluative or affective priming: Participants' reaction times to adjectives are faster

when those adjectives are preceded by similarly (versus dissimilarly) valenced primes.

According to this body of research, an *automatic attitude* consists of evaluative information that is associated with the attitude object in memory (Fazio, 2000; see also Ferguson & Bargh, in press). This interpretation presupposes so-called symbolic, localist models of memory wherein representations of semantically related objects are interconnected within an associative network (e.g., Anderson, 1983; Anderson & Bower, 1973; Collins & Loftus, 1975; Smith, 1996). Upon perception of an attitude object, activation automatically spreads from the object representation to the evaluation of the object, and then possibly to evaluatively similar objects (see Fazio, 2000; Ferguson & Bargh, in press).

Researchers have asserted that the activation of the evaluative information in memory is automatic for two main reasons. The first is that the stimulus-onset asynchrony (SOA) between the prime word and target word in evaluative priming paradigms is too brief a delay to allow strategic (i.e., nonautomatic) responding to occur (e.g., Fazio et al., 1986; cf. Klauer, Rossmagel, & Musch, 1997; Neely, 1976, 1977). The second reason is that according to evaluative priming research, attitudes toward objects can be activated even when the attitude objects are subliminally presented, thereby demonstrating automatic activation (Greenwald, Draine, & Abrams, 1996; Greenwald et al., 1989; Wittenbrink et al., 1997).

Although participants are typically aware of the presentation of the prime words in the sequential priming paradigm, they are unaware that their automatic reactions to the prime words are being measured. This is in clear contrast to explicit attitude measurement, wherein participants are pointedly asked to report their attitudes, which requires them first to estimate their attitude, and then decide whether and how to report that attitude. Because of the contrast in the way in which implicit and explicit attitudes are measured, and the consequent difference in participant's awareness of what is being measured, researchers are systematically exploring what exactly is captured by implicit measurement as compared with explicit measurement. In other words, in what ways, if any, do automatic attitudes (i.e., those measured implicitly) differ from strategic attitudes (i.e., those measured explicitly)?

For example, researchers have studied the extent to which automatic attitudes predict subsequent attitude-relevant behavior (e.g., Fazio et al., 1995; Swanson, Rudman, & Greenwald, 2001), whether automatic attitudes correspond to strategic attitudes (e.g., Cunningham et al., 2001; McConnell & Leibold, 2001), and how automatic attitudes differ from strategic attitudes in terms of stability across time and sus-

ceptibility to persuasion and change (e.g., Wilson et al., 2000). The extent to which the traditional conceptualization of attitudes can be applied to implicit attitudes has also been discussed (see Banaji, 2001).

Although there are dimensions along which strategically recalled attitudes differ from implicitly measured attitudes, such as perhaps the immediacy with which the attitudes are invoked, we focus on the possibility that the two types of measurement capture attitudes that are similarly functional for the perceiver. Specifically, given that both strategic and automatic attitudes provide perceivers with affective information regarding the respective objects, we apply the functional perspective on strategic attitudes to those attitudes that are automatically activated.

THE KNOWLEDGE FUNCTION OF AUTOMATIC ATTITUDES: OVERVIEW

Many researchers and theorists have argued that an important function of both strategic and automatic attitudes is the delivery of affective information about the attitude object (Fazio, 1989). Katz (1960) and Smith et al. (1956) referred to such appraisal as the knowledge function and identified it as one of the four primary functions of attitudes. Fazio (1989) claimed that this function is the most important because it refers to all attitudes, regardless of the intensity or direction of the attitude. He states that "regardless of why the individual's attitude took on a particular valence, the mere possession of any attitude is useful to the individual in terms of orienting him or her to the object in question" (p. 172).

This behavioral orientation toward the object is an integral reason that an attitude is conceptualized as functional (e.g., Fazio, 1989; Lewin, 1935; Pratkanis et al., 1989). Attitudes provide evaluative information about the respective objects that then helps orient the person toward the objects in an appropriate manner, whereby the person can avoid threatening or unpleasant objects and approach safe or pleasing objects. Automatic attitudes are thus functional because they can quickly and effortlessly provide vital information to the perceiver, saving him or her the trouble of having to repeatedly and consciously figure out whether an object is safe or dangerous, pleasing or displeasing (Smith et al., 1956, p. 41).

Yet although researchers have discussed the functional nature of automatic attitudes (Fazio, 1989), previous literature suggests that the functional nature of automatic attitudes may be constrained relative to that of strategic attitudes. We discuss two proposed limitations to the function of automatic attitudes. The first limitation concerns the de-

gree to which automatic attitudes are sensitive to the context within which attitude objects are perceived. Specifically, strategic attitudes are widely acknowledged to be context sensitive: A strategic attitude depends on the context in which the object is perceived, and consequently, the goals of the perceiver concerning that object (e.g., see Anderson, 1974; Bem, 1972; Fazio, 1987; Millar & Tesser, 1986; Olson, 1990; Schwarz & Clore, 1983; Strack, 1992; Tesser, 1978; Wilson, Dunn, Kraft, & Lisle, 1989; Wilson & Hodges, 1992). Such sensitivity of explicit attitudes is clearly important because it allows the person to behave in a situation-relevant fashion toward the object, according to his or her goals within the situation.

In contrast, automatic attitudes have sometimes been conceptualized as context independent (e.g., Fazio et al., 1995; Wilson & Hodges, 1992; Wilson et al., 2000), with some recent exceptions (e.g., Dasgupta & Greenwald, 2001; Wittenbrink, Judd, & Park, 2001). Such insensitivity would undermine the functional nature of automatic attitudes because automatic attitudes would be unable to account for situation-specific constraints regarding the object. We address this issue by exploring the sensitivity of automatic attitudes—if the functional nature of automatic attitudes rests on the delivery of information about how to relate to the object (e.g., approach or retreat), the information should be sensitive to the meaning of the object according to different goals of the perceiver. That is, an automatic attitude toward a given object should vary according to the utility of the object for the person, which may change across contexts and time. We review recent research that addresses this issue and suggests that automatic attitudes are contextually sensitive (Ferguson & Bargh, 2002).

The second limitation concerns the extent to which automatic attitudes are flexible to novel attitude objects, or whether they depend solely on preexisting, stored representations of evaluation. Research shows that people can strategically evaluate novel objects and can integrate evaluative information from a variety of different sources (e.g., Anderson & Rosenfeld, 1988; Bechtel & Abrahamson, 1991; Carlston and Smith, 1996; Fiedler, 1996; Smith, 1996; Smith & DeCoster, 1999). Although work outside the area of attitudes suggests that people can automatically integrate evaluative information (e.g., Fiedler, 1996), the literature on automatic attitudes instead suggests that automatic activation of an attitude depends solely on a stored, preexisting representation of an evaluation, which is associated with the object in memory (Fazio, 2000; Ferguson & Bargh, in press; Wilson et al., 2000). This view suggests that automatic attitudes can be invoked only in response to previously appraised objects.

Furthermore, it is presumed that the construction of an attitude

for a novel object requires some degree of conscious effort and thus cannot occur automatically (Fiske & Pavelchak, 1986; Wilson & Hodges, 1992; Wilson et al., 2000). This research suggests that people would be unable to automatically evaluate a novel object because of their inability to integrate multiple sources of evaluative information. To address this issue, we describe recent research that demonstrates that automatic attitudes are constructed in response to novel stimuli, thereby suggesting that automatic attitudes are flexible and constructive (Duckworth, Bargh, Garcia, & Chaiken, in press; Ferguson & Bargh, in press).

SENSITIVITY OF AUTOMATIC ATTITUDES

The primary function of an automatic attitude is that it quickly and effortlessly delivers information about how to relate to the attitude object (Fazio, 1989). For example, an object that evokes negativity can be immediately avoided or at least closely monitored, and an object that provokes positivity can be approached (e.g., Lewin, 1935; Rosenberg, 1956). This evaluative information is essential for how the perceiver understands the environment and structures his or her actions toward a given object (Fazio, 1989). Thus, the automatic evaluation of an object is functional because it guides behavior toward that object.

It is important to note that much theory asserts that behavior toward a given object will depend on the meaningfulness of that object within the situation in which it is encountered. That is, the meaning of an object is determined by the relation of the object to the person's goals within that situation. This perspective suggests that knowledge about objects is context bound (e.g., Fiske, 1992; Glenberg, 1997; Lewin, 1935). For example, Glenberg (1997, p. 5) argues that "to a particular person, the *meaning* of an object, event, or sentence is what that person can do with the object, event, or sentence" (*italics added*). For example, "depending on the context, a [C]oke bottle can be used to quench thirst, as a weapon, a doorstop, or a vase. That is, *meaning* depends on the context" (Glenberg, 1997, p. 6, *italics added*).

From this perspective, therefore, the evaluative knowledge about how to relate to an object must be sensitive to one's actual goals concerning that object, which can vary across contexts. If one is thirsty, a Coke bottle can be used to quench thirst and should therefore be evaluated positively. If, on the other hand, one is confronted by a mugger who is holding a Coke bottle as a weapon, the bottle should be evaluated negatively. Lewin (1935, p. 78) presaged this argument by asserting that "the valence of an object usually derives from the fact that the object is a means to the satisfaction of a need, or has indirectly some-

thing to do with the satisfaction of a need." The evaluative information of an object should therefore depend on the object's usefulness for the perceiver's goal(s). When an object should be approached, the attitude should be positive, and when the object should be avoided, the attitude should be negative. As Lewin (1935, p. 81) states, "the kind (sign) and strength of the valence of an object or event thus depends directly on the momentary condition of the needs of the individual concerned."

This perspective has been readily applied to explicit attitudes but not, as yet, to automatic attitudes (see Ferguson & Bargh, 2002). In fact, though some recent research has shown that automatic attitudes can vary if different information about the object is made salient (e.g., Dasgupta & Greenwald, 2001; Wittenbrink et al., 2001), other research suggests that automatic attitudes are contextually independent and rigid—largely impervious to change or fluctuation (Wilson & Hodges, 1992; Wilson et al., 2000; cf. Nelson & Bouton, Chapter 3, this volume). To examine the degree to which automatic attitudes are functional in the sense that they provide contextually sensitive information about how to relate to an object, we review some recent research that addresses this issue. We first explore the reasons that automatically activated evaluative knowledge should be context sensitive and then review a series of three studies that demonstrate such sensitivity.

Why Should Automatic Attitudes Be Sensitive to Context?

Much research has demonstrated that explicit attitudes vary according to contextual factors such as mood, previously activated attitude-relevant information, and experimental expectations, for instance (e.g., see Anderson, 1974; Bem, 1972; Fazio, 1987; Millar & Tesser, 1986; Olson, 1990; Schwarz & Clore, 1983; Strack, 1992; Tesser, 1978; Wilson et al., 1989; Wilson & Hodges, 1992). As discussed in the introduction of this chapter, when participants want to please the experimenter and are able to guess how the experimenter wants them to respond, they often respond accordingly (e.g., Orne, 1962; Rosenthal & Jacobson, 1966). Moreover, participants will report more egalitarian attitudes toward marginalized group members if they believe that they will be held accountable or if they want to display a fair and just way of thinking (e.g., Gaes et al., 1978; Ostrom, 1973).

From this body of research, it is evident that when people explicitly and purposefully report their attitudes, the direction and intensity of their attitudes depend on a variety of factors, perhaps including their underlying, unadulterated actual attitudes. The regular fluctuation of explicitly measured attitudes have led some researchers to argue that

stored attitudes do not exist, but instead are always constructed on the spot, according to temporarily activated subsets of attitude-relevant information in memory (e.g., Anderson, 1974; Tesser, 1978).

There is no a priori reason to assume that automatically activated evaluative information would not also be sensitive to the context in which the attitude object is measured. Many theorists claim that attitude objects are associated with a complex array of memories, including exemplars (e.g., Abelson, 1976, 1981; Bower, 1981; Eagly & Chaiken, 1993; Fishbein & Ajzen, 1975; Fiske & Pavelchak, 1986; Schank & Abelson, 1977). Given such an array of memories associated with a given attitude object, it seems likely that some object-relevant memories may be positively valenced and others may be negatively valenced (see Dasgupta & Greenwald, 2001). For example, someone may have many positive memories of drinking a sugary soda but may also have negative memories of the harmful effects of ingesting large quantities of caffeine. The context in which the soda product is perceived may determine the nature of the automatic attitude toward the drink, just as the context may influence the explicitly reported attitude.

Recent Research

Three experiments were conducted to test the sensitivity of automatic attitudes. The first two experiments explored the degree to which the framing of attitude objects influences the automatic attitudes toward those objects. Attitude objects that have been used in previous research and are considered to be fairly stable across time were included as the attitude objects of interest (e.g., see Bargh et al., 1992). We included normatively positive and negative attitude objects (e.g., dentist, chocolate). The methodology was largely the same for the two experiments. On each given trial, a prime word that represented an attitude object (i.e., "dentist") was presented for 250 milliseconds and was then followed by the presentation of an adjective (e.g., "generous," "awful"). Participants were asked to rate the adjectives as "good" or "bad" as quickly and accurately as possible.

In both experiments, we were interested in a specific set of attitude objects and wanted to demonstrate that each attitude object of interest could be framed as something safe and pleasing versus dangerous and displeasing, and that the automatic attitude toward that object would reflect the framing. Accordingly, each trial that contained an attitude object of interest (as the prime) was yoked with a preceding trial in which the prime word was semantically related to the subsequent attitude object of interest. Thus, across all trials, we were concerned with only those trials that contained attitude objects of interest. The rest of

the trials simply served to *frame* the way in which the attitude objects of interest were perceived.

The critical between-subject manipulation was whether the attitude objects of interest were preceded by attitude objects (as the primes in the preceding trials) that *matched* versus *mismatched* in valence with the traditional valence of the attitude objects of interest. For example, those participants who were in the condition in which the attitude objects were always framed according to their traditional valence would first see a trial with the attitude object *drill* (followed by an adjective), and then in the following trial see the attitude object *dentist* (followed by an adjective). This set of two trials for the attitude object *dentist* would be repeated with multiple positive and negative adjectives. We expected participant's reaction times to the negative adjectives (those that followed "dentist") to be shorter than their reaction times to the positive adjectives (those that followed "dentist") because we expected the word "drill" to frame *dentist* as something painful and negative.

In contrast, those participants who were in the condition in which the attitude objects were framed in an opposite manner to their traditional valence would first see "doctor" (followed by an adjective) and then, in the next trial, see "dentist" (followed by an adjective). In this case, we expected participants' reaction times to the positive adjectives (those that followed "dentist") to be shorter than their reaction times to the negative adjectives (those that followed "doctor") because we expected the word "doctor" to frame the attitude object *dentist* as something safe, helpful, and positive. The critical analysis across all attitude objects of interest was whether the valence of an attitude changed depending on the previously presented, semantically related information in the preceding trial.

In addition, the first experiment also explored the stability of automatic attitudes across two measurement points. Although previous research has demonstrated that automatic attitudes might be unreliable across time (e.g., Cameron, Alvarez, & Bargh, 2000), we hypothesized that such automatic appraisals should be stable, as long as the context within which the attitude object is perceived remains stable. Thus, participants completed the sequential priming paradigm twice in the same experiment. After they completed the paradigm the first time, they completed a number of control measures for approximately half an hour. They then completed the same exact priming paradigm again. We predicted that if the context determines the automatic attitude toward a given attitude object, there should be high and significant correlations between the automatic attitudes in the first measurement and those in the second measurement.

The hypothesis was supported by the pattern of data from each experiment. In the condition in which attitude objects in the preceding trials were evaluatively consistent with the traditional valence of the attitude objects of interest, automatic attitudes toward those attitude objects of interest were consistent with traditional conceptions. The word "dentist" automatically evoked a negative attitude, for instance. Comparatively, in the condition in which the previously presented attitude objects were evaluatively inconsistent with the traditional valence of the attitude objects of interest, automatic attitudes toward those attitude objects were reliably and strongly reversed as compared with the other condition. For these participants, the word "dentist" automatically evoked a positive attitude.

Finally, the first experiment also demonstrated the reliability of the automatic attitude measurement. We examined the reliability by comparing the correlations between reaction times to positive and negative adjectives as a function of the preceding attitude object, across the two time measurements. The correlations were highly positive and significant, thereby suggesting that if the context within which an attitude object is perceived remains constant, the nature of the automatic attitude is reliable. In addition, the automatic attitudes toward the objects of interest in the first measurement significantly predicted the automatic attitudes in the second measurement.

These first two experiments demonstrated that automatically activated attitudes can be completely reversed by manipulating the nature of recently activated, semantically related (to the attitude object) information. When the attitude object was framed as a dangerous or displeasing object (e.g., *drill* highlights the danger of *dentist*), the automatic attitude toward that object was negative. In contrast, when the object was framed as a safe or pleasing object (e.g., *doctor* emphasizes the safety, competence, and prestige of *dentist*), the automatic attitude toward that object was positive. These two experiments demonstrate that the evaluative knowledge that is delivered about an object is sensitive to the meaningfulness of the object to the perceiver. When negative aspects of the object are salient, the evaluative information is negative. When positive aspects of the object are highlighted, the evaluative information is positive. In addition, when the context is held constant, automatic attitudes are highly reliable.

A third experiment was conducted to further examine the degree to which the knowledge that an automatic attitude delivers is sensitive to the goals of the perceiver. Although the previous two experiments suggested that the evaluative information that is automatically activated about an object depends on which aspect of the object is made salient, we wanted to test more directly whether a current activated goal would

render some attitude objects more useful than when the goal is not activated. In addition, we were interested in the extent to which automatic attitudes would be responsive to a currently held goal, versus a goal that had already been completed.

In the third experiment, we focused on the goal of achievement. Participants sat in individual cubicles and were told that they were going to play a game in which they would be asked to make words from 15 Scrabble letter tiles. They were all told that they would have 5 minutes to make different words from the letter tiles. There were two between-subject variables. The first concerned whether an achievement goal was induced. Some participants were told that the game measured their verbal skills and their potential to succeed in various academic environments. These participants were also told that their performance would be compared to the performance of other students and so they should do as well as possible.

The same participants were told that they would receive points for their performance. Specifically, they were told that they would receive 1 point for every word, 5 points for every noun, and 7 points for every noun that started with the letter "c." They were instructed to achieve as many points as possible and to write down each word on a separate sheet of paper. The participants were timed, and the experimenter wore a stopwatch that she used in front of them to emphasize the need to work quickly.

The other participants were simply told to make words out of the letters for 5 minutes and that they would be asked questions about the experience of making words after the game was over. They were told that they did not need to write anything down, remember anything, or show the experimenter any of the words, because they would be asked questions later. Nothing about points was mentioned.

After all participants played the game for 5 minutes, they then completed a computer task that consisted of a sequential priming paradigm in which their automatic attitudes toward several attitude objects were measured. These words represented attitude objects that were useful for those who were trying to achieve (e.g., "words," "nouns," "c," "points," "compete").

The second independent variable concerned whether participants were told that they would play a second round of the game after this computer task, or that they were finished with the Scrabble game. Some participants were told that there would be two rounds of the game and that they would play the first 5-minute round and then complete a task that would clear their minds before going onto the second round of the game. The other participants were told that they would play the Scrabble game for 5 minutes, and would then start a computer task. Thus,

during the priming paradigm, half of the participants believed that they would play the Scrabble game again in a couple of minutes and the other half believed that they were finished with the game.

The attitude objects presented in the priming paradigm were useful for those participants who were given an achievement goal. We therefore expected that their automatic attitudes toward these attitude objects would be more positive than those of participants for whom these attitude objects were not useful. However, we also predicted that those in the achievement goal condition who expected to play again would demonstrate the most positive attitudes, for these participants still held the goal of performing well in the game. In contrast, those in the achievement condition who thought that they were finished with the game would not evaluate the objects as positively because the objects would no longer be useful for them at that point.

The results supported this hypothesis. According to an analysis of the reaction times to the positive and negative adjectives, automatic attitudes for the objects were reliably more positive for those in the achievement condition who thought that they would be playing again in several minutes, as compared with those of participants in the other three groups. That is, for participants in the achievement condition who thought they would play again, reaction times to the positive adjectives were significantly shorter than reaction times to the negative adjectives. For those in the other three conditions, automatic attitudes toward the objects were marginally negative, as indicated by shorter reaction times to negative versus positive adjectives.

It should be noted that the performance of those in the achievement condition and the two-rounds condition did not differ from the performance of those in the achievement and the one-round condition. An important point, then, is that the interaction described earlier suggests that the positive attitudes toward the objects did not result from a learning mechanism wherein those who performed better with regard to the objects generated more positivity toward the objects. Rather, the findings demonstrate that only those for whom the objects were still useful automatically evaluated the objects as positive. Those for whom the objects were no longer useful evaluated the objects in a neutral to negative manner.

In sum, this recent research illustrates two aspects of the knowledge function of automatic attitudes. The first is that, in general, an automatic attitude toward an object is sensitive to the current utility of that object, according to the context in which it is perceived. For example, if an object is framed as painful or displeasing (e.g., when the word "drill" precedes the attitude object "dentist"), the automatic attitude is negative. The third experiment further suggests that automatic attitudes

are also sensitive to whether a goal is currently in place or has just been completed. These findings indicate that the knowledge provided by automatic attitudes is sensitive to the context in which the respective objects are perceived.

Finally, an important result is that this third experiment also suggests that automatic appraisal does not depend on highly habitual, familiar responses to the attitude object and can instead reflect very recently learned information about that object, which is contrary to previous research (e.g., Bargh, 1989, 1997; Smith & Lerner, 1986). Participants were able to automatically evaluate attitude objects as positive, even when the utility of those objects was learned only 5 minutes prior to the measurement of the attitudes toward those objects.

FLEXIBILITY OF AUTOMATIC ATTITUDES

We now turn to the second dimension concerning the knowledge function of automatic attitudes—namely, the degree to which automatic attitudes are flexible in response to novel information or objects. If the function of automatic attitudes is to deliver evaluative information quickly and effortlessly in order to aid the perceiver in his or her behavior toward the objects as well as his or her understanding of the situation, then it would be an advantage to be able to appraise novel objects. In other words, the ability to automatically integrate evaluative information from individual features of a novel object in order to provide a summary evaluative response to the object, instead of relying solely on preexisting stored representations, seems like an advantage. This would mean that an integration of evaluative information concerning novel objects could occur immediately and without the necessity of conscious attention.

Previous research suggests that automatic attitudes depend on previously stored evaluative associations (e.g., Bargh et al., 1992; Bargh et al., 1996; Fazio et al., 1986; Fiske & Pavelchak, 1986). For example, Fazio et al. (1986) refer to the association between the representation of the attitude object and the representation of the evaluation of that object in their discussion of attitude activation. On this subject, they state, "just as a knowledge structure concerning some object may consist of bits of information organized in a network of associations to the object, so too may affect be linked to the object. Furthermore, just as activation can spread from one node in the network to another, the present data indicate a spontaneous spreading of activation from the object to the affective association" (p. 236).

In addition, theorists have argued that an automatic activation of more than one evaluative representation requires conscious effort and cannot be accomplished automatically (Fiske & Pavelchak, 1986; Wilson & Hodges, 1992; Wilson et al., 2000). For example, Wilson et al. (2000) argue that some attitudes are stored in memory and can therefore be automatically activated, whereas other attitudes are not stored in memory and require an on-line integration of evaluative information from numerous sources. This integration requires some degree of strategic effort (i.e., is not automatic). According to this perspective, people are able to evaluate novel objects only if they can effortfully integrate information about the features of the object.

No previous research has directly examined the extent to which people are able to automatically evaluate novel objects. For example, although it may seem as though this topic has been addressed by the work of Zajonc and colleagues on mere exposure (e.g., Monahan, Murphy, & Zajonc, 2000; Murphy & Zajonc, 1993), this research has actually shown only that the explicit liking of stimuli increases as the exposure of the stimuli increases. Although this research demonstrates how the frequency of exposure influences a person's explicit positivity toward the stimuli, it does not address whether and how people are automatically evaluating each novel stimulus.

Recent Research

Across four experiments, Duckworth et al. (in press) found evidence that participants were able to nonconsciously and automatically evaluate novel auditory and visual stimuli. In the first two experiments, participants were asked to pronounce a series of adjectives that appeared one at a time in the center of a computer screen. Each adjective was preceded by an auditorily presented attitude object that was either a positive or negative word or nonsense (novel) word. The words were nouns that had been used in previous automatic evaluation experiments (e.g., "cancer," "music"; see Bargh et al., 1992). The nonsense (novel) words were composed of two-syllable utterances and were unfamiliar to the participants. The valence of the novel words was determined by a separate pilot study in which participants explicitly evaluated the novel stimuli (as was the case across all four experiments).

The pattern of data from each experiment demonstrates the signature evaluative priming effect, both when the attitude objects were familiar and when they were novel. Specifically, participants were faster at pronouncing the target adjectives when the preceding attitude objects were of the same (versus opposite) valence, regardless of whether the

CONCLUSION

attitude objects were novel. The results from these two experiments show that people can automatically evaluate novel sounds as good or bad without any corresponding representation of the meaning of those novel sounds in memory.

In the third experiment, Duckworth et al. (in press) used pictorial stimuli as the attitude objects, including pictures of both familiar and novel attitude objects; the novel pictures were abstract art images with which participants were unfamiliar. As predicted, participants pronounced the target adjectives more quickly when the adjectives were preceded by attitude objects of the same (vs. opposite) valence, again, irrespective of the novelty of the pictures. The findings replicated the results from the first two experiments, and suggest that participants are able to automatically appraise novel images as positive or negative.

The fourth experiment addressed the degree to which the findings would generalize to a different response task, rather than pronunciation. Participants were asked to make either approach motions by pulling a lever toward them or avoidance motions by pushing a lever away from them (see Chen & Bargh, 1999; Wentura, Rothermund, & Bak, 2000). Novel abstract art images were used as prime stimuli, as in the third experiment, and appeared one at a time on a computer screen. Participants were told to either push or pull the lever in response to each image. It was expected that if participants were able to automatically evaluate the novel objects, positive evaluations should facilitate approach arm motions, as compared with negative evaluations, and negative evaluations should facilitate avoidance arm motions relative to positive evaluations. The results support this pattern, thereby showing that people are able to automatically evaluate novel stimuli and that this automatic appraisal influences the speed with which participants can behave toward the prime stimuli.

Together, these four experiments suggest that automatic attitudes are flexible in that they can be constructed in response to objects with which participants are unfamiliar. This finding suggests that the automatic appraisal of objects does not require preexisting, stored attitude representations associated with those objects. In addition, the findings suggest that constructive processes in appraisal do not require effortful processing, contrary to some recent theorizing (Wilson & Hodges, 1992; Wilson et al., 2000).

Most important, these experiments suggest a parallel between strategic attitudes and automatic attitudes. Both types of attitudes can refer to novel stimuli, presumably by integrating evaluative information from multiple sources in order to deliver an evaluation about the object as a whole.

This chapter has considered the implicit measurement of people's attitudes and the extent to which such automatic attitudes are functional. Researchers have asserted that the delivery of evaluative knowledge about an object is the primary function of attitudes, and that automatic attitudes accomplish this as well and are therefore functional. However, previous research suggests a limit to the extent to which automatic attitudes are functional, relative to strategic attitudes.

In particular, although strategic attitudes are thought to be sensitive to the context in which the attitude object is perceived and measured, automatic attitudes have been conceptualized as stable and impervious to temporary influences of attitude-irrelevant factors (Fazio et al., 1995; Wilson & Hodges, 1992; Wilson et al., 2000). Indeed, the purported stability of implicit attitude measurement is one of the reasons researchers began to prefer it over explicit measurement (e.g., see Banaji, 2001). However, in line with recent research suggesting that automatic attitudes may depend on the attitude-relevant exemplars to which the perceiver is exposed (e.g., Dasgupta & Greenwald, 2001), we reviewed evidence from three experiments that suggest that automatic attitudes are sensitive to the context in which the object is perceived, and consequently, to the current goals of the perceiver. For example, participant's automatic attitudes toward objects depended on whether the attitude objects were framed as safe and pleasing or as dangerous and displeasing. Further, participants' automatic attitudes toward a group of objects were positive when those objects were useful to the participants' goals.

Although strategic attitudes have been conceptualized as flexible and sometimes constructive, automatic attitudes have been presumed to be stable across time and rather inflexible (e.g., Wilson et al., 2000). Moreover, research and theory have suggested that whereas people can strategically evaluate novel objects, they are unable to do so automatically (e.g., Fiske & Pavelchak, 1986; Wilson & Hodges, 1992). In contrast, we reviewed a series of four experiments that demonstrate that people do immediately and automatically appraise objects with which they are unfamiliar. This finding was replicated with two different types of novel stimuli (audibly presented "words," and images) and across two different response tasks (pronunciation and arm movements).

These two sets of studies suggest that automatic attitudes are functional in a sensitive and flexible way. They provide the perceiver with evaluative information about an object in a way that is sensitive to the meaningfulness of the object within the current situation and relative to the person's goals. Further, people seem able to automatically ap-

praise novel objects (i.e., those with no corresponding representation in memory). Together, these findings suggest that automatic attitudes can be considered functional to a similar extent as attitudes that are strategically recalled.

REFERENCES

- Abelson, R. P. (1976). Script processing in attitude formation and decision making. In J. S. Carroll & J. W. Payne (Eds.), *Cognition and social behavior* (pp. 33-45). Hillsdale, NJ: Erlbaum.
- Abelson, R. P. (1981). Psychological status of the script concept. *American Psychologist*, *36*, 715-729.
- Ajzen, I., & Sexton, J. (1999). Depth of processing, belief congruence, and attitude-behavior correspondence. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 117-138). New York: Guilford Press.
- Allport, G. W. (1935). Attitudes. In C. Murchison (Ed.), *Handbook of social psychology* (pp. 798-844). Worcester, MA: Clark University Press.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, DC: Winston.
- Anderson, J. A., & Rosenfeld, E. (1988). *Neurocomputing: Foundations of research*. Cambridge, MA: MIT Press.
- Anderson, N. H. (1974). Cognitive algebra: Integration theory applied to social attribution. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 7, pp. 1-101). New York: Academic Press.
- Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, J. S. Nairne, I. Neither, & A. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117-150). Washington, DC: American Psychological Association.
- Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 3-51). New York: Guilford Press.
- Bargh, J. A. (1990). Auto-motives: Preconscious determinants of social interaction. In E. T. Higgins & R. M. Sorrentino (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (Vol. 2, pp. 93-130). New York: Guilford Press.
- Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 10, pp. 1-61). Mahwah, NJ: Erlbaum.
- Bargh, J. A., Chaiken, S., Gowerder, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology*, *62*, 893-912.
- Bargh, J. A., Chaiken, S., Raymond, P., & Hynes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology*, *32*, 104-128.
- Bechtel, W., & Abrahamson, A. (1991). *Connectionism and the mind: An introduction to parallel processing in networks*. Oxford, UK: Basil Blackwell.
- Bem, D. J. (1972). Self-perception theory. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 6, pp. 1-62). New York: Academic Press.
- Bower, G. H. (1981). Mood and memory. *American Psychologist*, *36*, 129-148.
- Cameron, J. A., Alvarez, J. M., & Bargh, J. A. (2000). *Examining the validity of implicit and explicit measures of prejudice: Is there really a bona fide pipeline?* Paper presented at the 1st annual meeting of the Society for Personality and Social Psychology, Nashville, TN.
- Carlston, D. E., & Smith, E. R. (1996). Principles of mental representation. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 184-210). New York: Guilford Press.
- Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach and avoid the stimulus. *Personality and Social Psychology Bulletin*, *25*, 215-224.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*, 407-428.
- Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measures: Consistency, stability, and convergent validity. *Psychological Science*, *12*, 163-170.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of liked and disliked individuals. *Journal of Personality and Social Psychology*, *81*, 800-814.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.
- Devine, P. G. (1995). Prejudice and out-group perception. In A. Tesser (Ed.), *Advanced social psychology* (pp. 467-524). New York: McGraw-Hill.
- Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology*, *33*, 510-540.
- Duckworth, K., Bargh, J. A., Garcia, M., & Chaiken, S. (in press). The automatic evaluation of novel stimuli. *Psychological Science*.
- Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich.
- Elms, A. C. (1975). The crisis of confidence in social psychology. *American Psychologist*, *30*, 967-976.
- Fazio, R. H. (1987). Self-perception theory: A current perspective. In M. P. Zanna, J. M. Olson, & C. P. Herman (Eds.), *Social influence: The Ontario Symposium* (Vol. 5, pp. 129-150). Hillsdale, NJ: Erlbaum.
- Fazio, R. H. (1989). On the power and functionality of attitudes: The role of attitude accessibility. In A. R. Pratkanis, S. J. Breckler, & A. G. Greenwald (Eds.), *Attitude structure and function* (pp. 153-180). Hillsdale, NJ: Erlbaum.
- Fazio, R. H. (2000). On the automatic activation of associated evaluations: An overview. *Cognition and Emotion*, *14*, 1-27.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013-1027.

- 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*, 299-298.
- Ferguson, M. J., & Bargh, J. A. (in press). The constructive nature of automatic evaluation. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion*. Hillsdale, NJ: Erlbaum.
- Ferguson, M. J., & Bargh, J. A. (2002). *Markers of meaningfulness: A motivational perspective of automatic attitudes*. Unpublished manuscript, New York University.
- Fiedler, K. (1996). Explaining and simulating judgment biases as an aggregation phenomenon in probabilistic, multiple-cue environments. *Psychological Review, 103*, 193-214.
- Fishbein, M., & Ajzen, I. (1974). Attitudes toward objects as predictors of single and multiple behavioral criteria. *Psychological Review, 81*, 59-74.
- Fishbein, M., & Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research*. Reading, MA: Addison-Wesley.
- Fiske, S. T. (1992). Thinking is for doing: Portraits of social cognition from da-guerreotype to laserphoto. *Journal of Personality and Social Psychology, 63*, 877-889.
- Fiske, S. T., & Pavelchak, M. A. (1986). Category-based versus piecemeal-based affective responses: Development in schema-triggered affect. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 167-203). New York: Guilford Press.
- Gaes, G. G., Kalle, R. J., & Tedeschi, J. T. (1978). Impression management in the forced compliance situation. *Journal of Experimental Social Psychology, 14*, 493-510.
- Glaser, J., & Banaji, M. R. (1999). When fair is foul and foul is fair: Reverse priming in automatic evaluation. *Journal of Personality and Social Psychology, 77*, 669-687.
- Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences, 20*, 1-55.
- Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three cognitive markers of unconscious semantic activation. *Science, 273*, 1699-1702.
- Greenwald, A. G., Klinger, M. R., & Liu, T. J. (1989). Unconscious processing of dichoptically masked words. *Memory and Cognition, 17*, 35-47.
- Greenwald, A. G., McChae, D. E., & Schwarz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464-1480.
- Himmelfarb, S. (1993). The measurement of attitudes. In A. H. Eagly & S. Chaiken (Eds.), *The psychology of attitudes* (pp. 23-84). Fort Worth, TX: Harcourt Brace Jovanovich.
- Jones, E. E., & Sigall, H. (1971). The bogus pipeline: A new paradigm for measuring affect and attitude. *Psychological Bulletin, 76*, 349-364.
- Katz, D. (1960). The functional approach to the study of attitudes. *Public Opinion Quarterly, 24*, 163-204.
- Katz, D., & Stotland, E. (1959). A preliminary statement to a theory of attitude structure and change. In S. Koch (Ed.), *Psychology: A study of a science* (Vol. 3, pp. 423-475). New York: McGraw-Hill.
- Kelman, H. C. (1974). Attitudes are alive and well and gainfully employed in the sphere of action. *American Psychologist, 29*, 310-324.
- Klauer, K. C. (1998). Affective priming. *European Review of Social Psychology, 8*, 63-107.
- Klauer, K. C., Rossnagel, C., & Musch, J. (1997). List context effects in evaluative priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 246-255.
- Lewin, K. (1935.). *A dynamic theory of personality*. New York: McGraw-Hill.
- McConahay, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91-125). San Diego, CA: Academic Press.
- McConnell, A. R., & Leibold, J. M. (2001). Relations among the Implicit Association Test, discriminatory behavior, and explicit measures of racial attitudes. *Journal of Experimental Social Psychology, 37*, 435-442.
- McGuire, W. J. (1986). The vicissitudes of attitudes and similar representational constructs in twentieth century psychology. *European Journal of Social Psychology, 16*, 89-130.
- Millar, M. G., & Tesser, A. (1986). Thought-induced attitude change: The effects of schema structure and commitment. *Journal of Personality and Social Psychology, 51*, 259-275.
- Monahan, J. L., Murphy, S. T., & Zajonc, R. B. (2000). Subliminal mere exposure: Specific, general, and diffuse effects. *Psychological Science, 11*, 462-466.
- Murphy, S. T., & Zajonc, R. B. (1993). Affect, cognition, and awareness: Affective priming with optimal and suboptimal stimulus exposures. *Journal of Personality and Social Psychology, 64*, 723-739.
- Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory and Cognition, 4*, 648-654.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General, 106*, 225-254.
- Nesdale, D., & Durkin, K. (1998). Stereotypes and attitudes: Implicit and explicit processes. In K. Kirchner & C. Spelman (Eds.), *Implicit and explicit mental processes* (pp. 219-232). Mahwah, NJ: Erlbaum.
- Olson, J. M. (1990). Self-inference processes in emotion. In J. M. Olson & M. P. Zanna (Eds.), *Self-inference processes: The Ontario Symposium* (Vol. 6, pp. 17-42). Hillsdale, NJ: Erlbaum.
- Orne, M. T. (1962). On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist, 17*, 776-783.
- Ostrom, T. M. (1973). The bogus pipeline: A new ignis fatuus? *Psychological Bulletin, 79*, 252-259.
- Pratkanis, A. R., Breckler, S. J., & Greenwald, A. G. (1989). *Attitude structure and function*. Hillsdale, NJ: Erlbaum.

- Rosenberg, M. J. (1956). Cognitive structure and attitudinal affect. *Journal of Abnormal and Social Psychology, 53*, 367-372.
- Rosenberg, M. J., & Hovland, C. I. (1960). Cognitive, affective, and behavioral components of attitudes. In C. I. Hovland & M. J. Rosenberg (Eds.), *Attitude organization and change: An analysis of consistency among attitude components* (pp. 1-14). New Haven, CT: Yale University Press.
- Rosenthal, R., & Jacobson, L. (1966). Teachers' expectancies: Determinants of pupils' IQ gains. *Psychological Reports, 19*, 115-118.
- Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition, 17*, 437-465.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.
- Schuman, H., & Johnson, M. P. (1976). Attitudes and behavior. *Annual Review of Sociology, 2*, 161-207.
- Schuman, H., & Kalton, G. (1985). Survey methods. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (3rd ed., Vol. 1, pp. 635-697). New York: Random House.
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgment of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology, 45*, 513-523.
- Sigall, H., & Page, R. (1971). Current stereotypes: A little fading, a little faking. *Journal of Personality and Social Psychology, 18*, 247-255.
- Smith, E. R. (1996). What do connectionism and social psychology offer each other? *Journal of Personality and Social Psychology, 70*, 893-912.
- Smith, E. R., & DeCoster, J. (1999). Associative and rule-based processing: A connectionist interpretation of dual-process models. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 323-336). New York: Guilford Press.
- Smith, E. R., & Lerner, M. (1986). Development of automatism of social judgments. *Journal of Personality and Social Psychology, 50*, 246-259.
- Smith, M. B., Bruner, J. S., & White, R. W. (1956). *Opinions and personality*. New York: Wiley.
- Strack, F. (1992). The different routes to social judgments: Experiential versus informational strategies. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 249-276). Hillsdale, NJ: Erlbaum.
- Swanson, J. F., Rudman, L. A., & Greenwald, A. G. (2001). Using the Implicit Association Test to investigate attitude-behavior consistency for stigmatized behaviour. *Social Cognition, 15*, 207-230.
- Tesser, A. (1978). Self-generated attitude change. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 11, pp. 289-338). New York: Academic Press.
- Thompson, M. M., Zanna, M. P., & Griffin, D. W. (1995). Let's not be indifferent about (attitudinal) ambivalence. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences*. Ohio State University series on attitudes and persuasion (Vol. 4, pp. 361-386). Hillsdale, NJ: Erlbaum.
- Thurstone, L. L. (1928). Attitudes can be measured. *American Journal of Sociology, 33*, 529-554.
- Warner, S. L. (1965). Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association, 60*, 63-69.
- Wentura, D., Rothermund, K., & Bak, P. (2000). Automatic vigilance: The attention grabbing power of approach- and avoidance-related social information. *Journal of Personality and Social Psychology, 78*, 1024-1037.
- Wicker, A. W. (1969). Attitude versus actions: The relationship of verbal and overt behavioral responses to attitude objects. *Journal of Social Issues, 25*(4), 41-78.
- Wilson, T. D., Dunn, D. S., Kraft, D., & Lisle, D. J. (1989). Introspection, attitude change, and attitude-behavior consistency: The disruptive effects of explaining why we feel the way we do. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 22, pp. 287-343). Orlando, FL: Academic Press.
- Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In A. Tesser & L. Martin (Eds.), *The construction of social judgment* (pp. 37-65). Hillsdale, NJ: Erlbaum.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107*, 101-126.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72*, 262-274.
- Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology, 81*, 815-827.