Beyond the attitude object:

Automatic attitudes spring from object-centered-contexts

Melissa J. Ferguson — **Cornell University**

John A. Bargh — **Yale University**

Invited chapter in B. Wittenbrink and N. Schwarz (Eds.), ***Implicit measures of attitudes*** **(pp. 216-246). NY: Guilford .**

Beyond the attitude object:

Automatic attitudes spring from object-centered-contexts

One of the most seminal developments in attitude research over the last two decades has been the advent of measures that capture people's attitudes in an implicit manner, as the current volume attests (see also Banaji, 2001; Bassili & Brown, 2005; Bargh, Chaiken, Govender, & Pratto, 1992; Blair, 2002; Fazio, 2001; Fazio, Sanbonmatsu, Powell, & Kardes, 1986; Fazio & Olson, 2003; Ferguson, in press; Greenwald & Banaji, 1995; Greenwald, McGhee, & Schwartz, 1998; Musch & Klauer, 2003). These measures gauge the degree to which people evaluate a given stimulus as good or bad without their awareness that their attitude toward the stimulus is being measured. Some implicit measures can capture a person's attitude toward a stimulus without the person even consciously registering the stimulus itself, much less that her or his attitude toward it is being assessed (e.g., Greenwald, Klinger, & Liu, 1989).

Implicit attitude measurement stands in contrast to the direct self-report methodology that has guided much of attitude research over the last century. Whereas implicit measures gauge attitudes covertly, self-report measures consist of directly asking people to report their attitude or evaluation of a certain stimulus (see Albarracín, Johnson, & Zanna, 2005; Eagly & Chaiken, 1993). For instance, the standard "feeling-thermometer" measure of attitudes requires respondents to report their warmth from 0 to 100 degrees toward some group or object. Even many of the subtler attitude measures developed in the late 1970's and 1980's required people to explicitly and deliberately report their attitude. The "bogus pipeline" developed by Jones and Sigall (1971), for example, was predicated on the assumption that people will intentionally report a more truthful attitude if they believe that any fabrication will be detected (see also McCauley & Stitt, 1978; McConahay, 1986). Because such measures allow the respondent to respond

strategically (and typically without time constraints), attitudes measured explicitly might reflect respondents' attempts at hiding their true feelings, or yielding to what they think the experimenter wants them to feel (e.g., see Orne, 1962; Rosenberg, 1969).

Contemporary implicit measures that gauge people's attitudes without their awareness thus represent a substantive departure from how attitudes have characteristically been measured over the last 100 years. Given this rather dramatic shift in measurement techniques, there has been a considerable amount of research activity directed toward understanding the precise nature of implicit attitudes.[1] Because explicit measures have been found to be highly sensitive to contextual factors at the time of measurement, a great deal of this research activity has examined how implicit attitudes might similarly dependent on the context (e.g., Barden, Maddux, Petty, & Brewer, 2004; Dasgupta & Greenwald, 2001; Ferguson & Bargh, 2004; Karpinski & Hilton, 2001; Livingston & Brewer, 2002; Lowery, Hardin, & Sinclair, 2001; Mitchell, Nosek, & Banaji, 2003; Wittenbrink , Judd, & Park, 2001). Additionally, differences between implicit and explicit attitudes have prompted the development (or modification) of theories concerning the representation and generation of evaluative information in memory more generally (e.g., Bassili & Brown, 2003; De Houwer, in press; Ferguson, in press; Gawronski, Strack, & Bodenhausen, in press; Mitchell et al., 2003).

Largely on the basis of this recent research and theory, we argue in the current chapter for a re-conceptualization of what automatic attitudes reflect. Rather than such an attitude reflecting the automatically activated evaluation of an "object," we argue that automatic attitudes would be more appropriately understood as automatically activated evaluations of "object-centered-contexts". An overt emphasis on the contextual dependence of such attitudes is reasonable, we argue, given the evidence accumulated over the last five years especially, and recent theoretical

frameworks concerning the generation of evaluations. Additionally, we also speculate that automatic attitudes may be especially reflective of contextual factors compared with more deliberate attitudes. In what follows, we review the contemporary definition of an automatic attitude, the research and theory suggesting their contextualized nature and relevant issues concerning attitude stability, and finally a consideration of why automatic attitudes may be more context-dependent than more deliberately generated attitudes.

What is an automatic attitude?

There are arguably three components of the definition of an automatic attitude. The first concerns the attitude itself. From the inception of the construct in the beginning of the 20th century (e.g., Albarracín et al., 2005; Eagly & Chaiken, 1993), attitudes have been understood as assessments of whether the corresponding objects are pleasing or displeasing. Although some definitions have included aspects of beliefs and behavior (e.g., Fishbein & Ajzen, 1975), they have always stressed the importance of the evaluative dimension, and contemporary definitions seem to rest primarily on the evaluative aspect. For example, one of the most currently widely accepted definitions stipulates that an attitude is a "psychological tendency that is expressed by evaluating a particular entity with some degree of favor or disfavor" (Eagly & Chaiken, 1993). Similarly, Fazio and colleagues (Fazio, 1986, 1995, 2001; Fazio et al., 1986) have argued that an attitude consists of the association in memory between an object and its (positive or negative) evaluation.

The second component refers to the nature of the attitude object. Traditionally, the *object* of an evaluation can consist of whatever the person is thinking about at the time of measurement. The only requisite for an "attitude object" (i.e., an object that is being evaluated) is that it is discriminable, or a subject of thought (Eagly & Chaiken, 1993). Attitude objects have most

often consisted of physical, concrete objects (e.g., apple, garbage), individuals or groups of people (e.g., Michael Jordan, Blacks), abstract ideals and values (e.g., equality, freedom), and issues and policies (e.g., abortion, voting). Within the automatic attitude literature, attitudes have most often been measured in response to concrete graspable objects, such as people and items (though see Ferguson, 2005). This literature suggests overall that an attitude is assumed to reflect one's evaluation of an object in isolation.

The third component of the definition of an automatic attitude of course entails the automatic way in which the attitude is generated. An attitude can be characterized as automatic when its activation does not require the perceiver's intention, awareness, or effort (see Bargh, 1994, 1997; Bargh et al., 1992). A large body of work has now demonstrated that attitudes can be activated under such circumstances, although the vast majority of the research in this area focuses on the *intentionality* dimension of automaticity. Explicit attitudes, in contrast, necessarily require the respondent's intention, awareness, and control. The two major types of attitude measurement currently used in the field thus differ in terms of the way in which evaluative information about an object is activated -- that the evaluation refers to a singular object is not questioned, and is not assumed to differ depending on the type of measurement.

The initial promise of automatic attitudes as context-dependent

The development of *implicit* measures was prompted by the accumulating, extensive evidence that *explicitly measured* attitudes are highly influenced by the context in which they are assessed (e.g., Anderson, 1974; Bem, 1972; Chaiken & Yates, 1985; Fazio, 1987; Feldman & Lynch, 1988; Forgas, 1992; Millar & Tesser, 1986; Olson, 1990; Schwarz & Clore, 1983; Schwarz & Bless, 1992; Schuman & Presser, 1981; Strack, 1992; Tesser, 1978; Tourangeau & Rasinski, 1988; Wilson, Dunn, Draft, & Lisle, 1989; Wilson & Hodges, 1992). Researchers

have demonstrated that people's explicit attitudes about the state of their marriage depend in part on whether it is sunny outside at the moment (Schwarz & Clore, 1983) and that a person's self-reported religiosity on a questionnaire depends in part on the nature of the preceding question (Salancik & Conway, 1975). Similarly, respondents who want to be perceived as egalitarian might intentionally report positive attitudes toward a stigmatized group to an experimenter, yet behave in a highly prejudiced manner in less public settings (e.g., LaPiere, 1934; McConahay, 1983).

Research demonstrating that explicit attitudes are influenced by observations of one's own behavior, current thoughts, mood, demand effects, and other contextual constraints led some researchers to suggest that explicit attitudes are highly contextually dependent, to the point that they are constructed on the spot (Anderson, 1974; Tesser, 1978). In other words, people might spontaneously construct their attitudes when they are asked to report them, rather than recalling and reporting preexisting attitudes in memory. Because of this constructive process at the time of measurement, contextual factors that might have little to do with the actual attitude itself nevertheless influence the response (see Schwarz & Bohner, 2001 for a review).

Given the inherent difficulties in interpreting data from explicit attitude measures, researchers began to search for less reactive, and therefore more effective, tools of attitude assessment. Implicit measures, which have been included in cognitive psychology methodologies for decades, were first employed to measure attitudes in 1986 by Fazio and colleagues (Fazio, Sanbonmatsu, Powell, & Kardes, 1986). Fazio et al. (1986) modified a sequential priming paradigm that was developed by Neely (1976, 1977). Whereas Neely used the priming paradigm to study the automatic activation of semantic information on perception of a stimulus, Fazio et al. (1986) modified the priming paradigm to examine the automatic

activation of solely evaluative information on perception of a stimulus. This paradigm has since been utilized as an implicit attitude measure (e.g., Fazio et al., 1995; Ferguson & Bargh, 2004; Wittenbrink et al., 1997, 2001), and comprehensive details of this measure are provided in Wittenbrink (this volume). Since the publication of the evaluative priming paradigm by Fazio and colleagues (Fazio et al., 1986), researchers have developed other means of implicitly assessing attitudes, including the Implicit Association Test (Greenwald et al., 1998; see also Lane, Banaji, Nosek, & Greenwald, this volume), various physiological tools and measures (see Ito & Cacioppo, this volume), and an assortment of other implicit, behavioral measures (De Houwer, 2003; De Houwer & Eelen, 1998; Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Koole, Dijksterhuis, & van Knippenberg, 2001; Nosek & Banaji, 2001; Payne, Cheng, Govorun, & Stewart, in press; von Hippel, Sekaquaptewa, & Vargas, 1997).

The development of implicit attitude measures was accompanied by the hope that such measures might prove less reactive and more covert than explicit attitude measures (Banaji, 2001; Fazio et al., 1995). From the beginning, automatic attitudes were assumed to be contextually independent (e.g., Bargh, 1999; Bargh et al., 1992; Bargh et al., 1996; Devine, 1989; Fazio et al, 1995; Wilson & Hodges, 1992; Wilson et al., 2000), to the point that an implicit attitude measure was regarded as a potential "bona fide pipeline" to people's inner attitudes (Fazio et al., 1995). One of the straightforward reasons for such an assumption is that implicit measures tap people's attitudes outside of their awareness, and thus without any strategic editing. That is, participants cannot possibly modify their responses if they are not aware of making such responses in the first place. Not only was it presumed that implicit measures are therefore impervious to strategic intentions, it was also assumed that they are able to capture a constant and unchanging summary evaluative index of the object (see Wilson et al., 2000).

Thus, at least initially, implicit measures seemed likely to provide clean, less variable indices of how people feel toward the stimuli in their environment. Researchers in attitude theory as well as many working on prejudice and stereotyping therefore initially embraced a view of implicit attitudes as context-independent and stable indices of people's "true" evaluations of objects, people, places, and concepts. Researchers studying socially sensitive topics such as prejudice toward various groups therefore began to rely on implicit attitude measures rather than explicit measures (e.g., Devine, 1989; Fazio et al., 1995; Swanson, Rudman, & Greenwald, 2001).

The contextual dependence of automatic attitudes

Over the last 5 years, in contrast to initial assumptions, there has been a wealth of evidence suggesting that automatic attitudes vary according to a host of factors at the time of measurement. This work suggests that one's implicit attitude toward a racial group, for example, will depend on one's recently activated memories about the group (Dasgupta & Greenwald, 2001), task instructions (Mitchell et al., 2003), the race of the experimenter (Lowery, Hardin, & Sinclair, 2001), one's active goals (Ferguson & Bargh, 2004; Sherman et al., 2003), one's chronic goals about egalitarianism (Barden et al., 2004; Maddux et al., 2005), and one's mood (Hermsen, Holland, & van Knippenberg, 2004). We review this evidence of contextual dependence by dividing the research into two broad groups – one group includes those experiments that directly constrained the types of object-relevant information accessible at the time of measurement, and another group that includes experiments that constrained the accessibility of object-relevant information more indirectly through goals and task instructions.

<u>Direct constraints (information about attitude object is manipulated physically in environment)</u>

The first batch of experiments to suggest the contextual dependence of automatic attitudes was published in 2001 in a special issue of the *Journal of Personality and Social Psychology* devoted to implicit measures of attitudes and stereotypes. In one of the articles, Dasgupta and Greenwald (2001) demonstrated that the way in which participants automatically evaluate a racial group depends on which group members are most accessible in memory at the time of measurement. Participants first looked at pictures of either liked White people and disliked Black people (pro-White group), or the reverse (pro-Black group). Participants then completed a race-based IAT (see Lane et al., this volume), and the results demonstrated that those in the pro-Black group displayed significantly less negative automatic attitudes toward Blacks than those in the pro-White group. Dasgupta and Greenwald showed that this effect also remained 24 hours later when participants came back to the lab to retake the race IAT. These findings reinforce the view that a given object is likely associated with a wide array of memories, some of which may differ in their evaluative connotation (e.g., Abelson, 1976, 1981; Barsalou, 1992; Bower, 1981; Fishbein & Ajzen, 1975; Fiske & Pavelchak, 1986; Schank & Abelson, 1977; Smith, 1992; Smith & Zarate, 1992). This means that the particular memories that are active at the time of an attitude measure will likely greatly influence the direction (positive, negative) of the attitude, a view that has also been applied to explicit attitudes (e.g., Schwarz & Bohner, 2001).

In that same volume, Wittenbrink et al. (2001) reported evidence that participants' automatic attitudes toward Black people depended on the context in which the people were encountered. When participants saw photos of a Black person in a picnic setting, for instance, they displayed significantly more positive automatic attitudes toward that person compared with those who saw photos of the same Black person standing by a wall with graffiti on it in an inner-

city scene. This work is consistent with that of Dasgupta and Greenwald (2001), and suggests that the attitudes that are activated on perception of a stimulus do not solely depend on the stimulus in isolation – they are heavily influenced by the rich context in which the stimulus is perceived.

Also in the same volume, Lowery et al. (2001) demonstrated that one way to alter the accessibility of object-relevant information is to vary the race of the experimenter. In this research, participants completed an implicit attitude measure on race, and the experimenter of the study was either Black or White. When the experimenter was Black, participants displayed significantly less negative automatic attitudes toward Blacks compared with when the experimenter was White. The authors interpreted these results as evidence that people's automatic attitudes are sensitive to the social pressures in a situation, such as the pressure to avoid being prejudiced in the presence of a person who belongs to a group that has traditionally been stigmatized. This suggests that interacting with a particular group member right before completing an implicit attitude measure alters one's automatic attitudes toward that group, just as the research by Dasgupta and Greenwald (2001) Wittenrbink et al. (2001) would suggest.

Karpinski and Hilton (2001) also investigated the influence of highly accessible, object-relevant information on automatic attitudes. In their research, however, they experimentally increased the accessibility of pro-elderly and anti-youth information through the repetitive viewing of pro-elderly and anti-youth stimuli. They found that those participants who had viewed the pro-elderly stimuli did in fact exhibit significantly more positive attitudes toward the elderly. This research is consistent with research on implicit stereotyping that shows that people's automatic stereotypes depend on the nature of recently activated, and repeatedly viewed

stereotype-relevant information (see also Blair, Ma, & Lenton, 2001; Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000; Rudman, Asmore, & Gary, 2001).

Other research has also demonstrated the potency of accessible object-relevant information on automatic attitudes toward the object. Livingston and Brewer (2002), for instance, tested whether the negative automatic evaluation of Black people was primarily due to facial features that were strongly prototypic of Blacks, versus the categorization of a person in terms of race or ethnicity. They found evidence for the former by showing that participants in fact displayed significantly more negative automatic attitudes toward those who possessed facial features that were prototypic of Black people, such as a wide nose and full lips, compared with less prototypic features. This suggests that certain group features are more strongly associated with negativity than others, and that making those features salient influences the type of attitude that is automatically generated toward the group. The degree of negativity of an individual's automatic attitude toward a given Black person, then, depends on the relative accessibility of prototypically Black versus non-prototypically Black features.

Indirect constraints (information about attitude object is manipulated via the perceiver's goals)

Research has also demonstrated that the accessibility of object-relevant information can be determined by the type of motivation or goal the person possesses at the time of measurement. Whereas the research reviewed earlier manipulated the accessibility of object information directly – that is, different kinds of information about the attitude object were included (or not) in participants' physical environment (e.g., Dasgupta & Greenwald, 2001; Wittenbrink, et al., 2001) – other research has presented the same physical stimulus but altered people's goal or task objective at the time of the measurement. For example, Mitchell et al. (2003) recently tested whether participants' automatic attitudes toward Blacks depend on the way in which they are

categorized. Participants completed an IAT, and were instructed to classify Black and White persons in terms of either career or race. Whereas the Black targets were all liked athletes, the White targets were all disliked politicians. The results demonstrated that when participants classified the targets in terms of career, they exhibited significantly more positive automatic attitudes toward Blacks than Whites. However, this pattern was reversed, and negativity toward Blacks was demonstrated, when participants classified the targets in terms of race. This research suggests that the categorization task determined the types of memories activated in response to each target person, and consequently the nature of the automatic evaluation toward that person. When Blacks were categorized in terms of their athleticism, more positive memories were presumably more activated and thus drove participants' automatic evaluative responses.

Maddux et al. (2005) recently demonstrated that people's automatic attitudes toward group members also depend on their chronic goals to avoid being prejudiced. They included participants who were low or high in the motivation to respond without prejudice, and asked them to complete an implicit attitude measure similar to the one developed and used by Wittenbrink et al. (2001). Participants saw photos of Black people who were in either negative (e.g., a prison) or positive (e.g., a church) contexts. The results showed that those who were low in the motivation to avoid prejudice demonstrated significantly more strongly negative automatic attitudes toward Black persons in the negative compared to the positive context. Interestingly, however, those who were high in the motivation to avoid prejudice actually showed weaker negative automatic attitudes in the negative compared to the positive context. For these participants, their negativity toward Blacks was apparently inhibited in response to a context that would normally provoke negativity. These findings suggest that people's chronic motivations

can determine the way in which they interpret, and evaluatively respond to the same physical stimulus (see also Moskowitz, Gollwitzer, Wasel, & Schaal, 1999).

Ferguson and Bargh (2004) demonstrated that people's currently active goals also influence the automatic evaluation of a range of (non-racial) stimuli, depending on the relation between participants' current goal and the utility of the stimuli. Participants who were pursuing a goal tended to automatically evaluate as positive those objects that could help them to achieve the goal, compared with other objects, and also those participants who had already reached the goal. For example, those participants who were playing a competitive word game exhibited significantly less negative automatic attitudes toward words that were relevant to the game (e.g., achieve, win, noun, c), compared with those who had finished the game only a minute earlier. This suggests that automatic attitudes can reflect the current or prospective utility of the corresponding objects, rather than solely the retrospective, or recent utility (see also Sherman et al., 2003). This work is in line with the classic notion that thinking (and evaluating) is in the service of action (e.g., Fiske, 1992; James, 1890).

In similar research, Sherman et al. (2003) tested whether participants' automatic attitudes toward smoking paraphernalia would depend on whether they had a current need to smoke. Participants in this study were all habitual smokers, and had either just smoked or not. Those who had <u>not</u> just smoked were found to automatically evaluate smoking-related objects as significantly more positive than those who had just fulfilled their nicotine needs. This work suggests, along with the work by Ferguson and Bargh (2004), that evaluations that are activated automatically reflect what the perceiver currently wants to do via-a-vis the objects, rather than solely what the perceiver has just recently done (see also Moors & De Houwer, 2001; Moors, De Houwer, & Eelen, 2004).

Theories on the generation of automatic attitudes

As the above review indicates, research over the last five years suggests that automatic attitudes are actually highly sensitive to the number of contextual factors in place at the time of measurement. These factors include the nature of recently encountered exemplars, the prototypicality of the attitude object itself, the context in which the target object is encountered, the task instructions, and the perceiver's chronic and temporary goals. Given this evidence, we now turn to theories of how attitudes are generated, and what they reflect. We first review contemporary theories, and then discuss a recent, more constructivist approach. Whereas the contemporary approaches assume that automatic attitudes refer to the respective objects in isolation, a more constructivist approach assumes that an automatic evaluation of an object will never be activated in isolation of other relevant contextual factors.

Automatic attitudes toward objects in isolation

There are currently two predominant theories about the representation and generation of automatic attitudes (e.g., Bargh et al., 1992; Fazio, 1986, 1995, 2001; Fazio et al., 1995; Fazio et al., 1986; Fiske & Pavelchak, 1986; Wilson et al., 2000). Whereas one can be understood as the "single-tag" perspective (e.g., Fazio, 1986, 1990, 2001), the other might be classified as the "dual-tag" perspective (Chaiken & Bargh 1993; Wilson et al., 2000). Both of these theories assume that evaluative information is associated in memory with an object, and importantly, that this is what is captured by implicit (and explicit) attitude measures. These two theories differ primarily in the number of summaries of evaluative information associated with the object, and the circumstances under which evaluative information might be activated.

Much of the research on the automatic activation of attitudes is based on the supposition that (many) attitude objects are each associated in memory with a positive or negative evaluative

index, or a summary evaluative "tag" (e.g., Bargh et al., 1992; Bargh et al., 1996; Fazio et al., 1986; Fiske & Pavelchak, 1986). After the activation of the object representation, the evaluative tag becomes activated and can then influence subsequent judgment and behavior. From the single-tag perspective, the strength of the association between an object and its evaluation is the main determinant of whether the attitude will be automatically activated in the first place, and then influence judgment and thought (Fazio, 1989, 1990; Fazio & Williams, 1986; Petty & Krosnick, 1995).

The single-tag view is largely based on a localist, symbolic model of memory that presupposes an associative network in which isolated nodes represent individual constructs, exemplars, or features of an object (e.g., Anderson & Bower, 1973; Collins & Loftus, 1975). These nodes are interconnected according to the amount of semantic relation between the nodes, with activation spreading along these links automatically upon perception of an object (e.g., Meyer & Schvaneveldt, 1971; Neely, 1976, 1977; Posner & Snyder, 1975; Shiffrin & Schneider, 1977). Such models have been referred to as "file-drawer" models of memory because the contents of memory remain inert and unchanged (as though stored in a file cabinet) until independent retrieval processes act on them (e.g., see Carlston & Smith, 1996; Smith, 1996).

According to the single-tag perspective, evidence of the contextual dependence of automatic attitudes might be explained though differences in categorization depending on the context. Because categories are likely associated with different summary evaluative tags (see Fiske & Pavelchak, 1986), differences in categorization might lead to differences in automatic attitudes. In this way, the single-tag perspective might explain at least some of the evidence for the contextual dependence of automatic attitudes by suggesting that the attitude objects were categorized in different ways across conditions. For example, from this perspective, one might

assume that in the work by Wittenbrink et al. (2001), the attitudes of those participants who saw photos of Black persons at a picnic reflected the evaluation of a subtype of Black persons (e.g., "non-city"). In contrast, the automatic attitudes of those who saw the photos of Black persons in inner-city scenes might have reflected a different, more negatively evaluated subtype (e.g., Weber & Crocker, 1983). In this way, even though there is variability in the attitudes depending on the context in which the perceiver encounters the object, each attitude is still assumed to reflect the evaluative tag associated with a single object (or category).

The single-tag perspective would have to address two issues in order to more fully explain the evidence for contextual dependence. The first is that some of the evidence suggests that the accessibility of positive information can be orthogonal to the accessibility of negative information relevant to an object (Cacioppo & Berntson, 1994; Cacioppo, Gardner, & Berntson, 1997). In particular, the findings from Ferguson and Bargh (2004) demonstrated that when participants were actively pursuing a competitive game, the accessibility of positive material associated with the relevant objects was unchanged, but the negativity associated with relevant objects was inhibited compared with those who had already finished the game. In another experiment, the automatic attitudes toward objects such as water and juice were compared across thirsty and non-thirsty people. The results showed that whereas the positivity associated with those objects was more accessible for thirsty (vs. non-thirsty) people, the accessibility of the negativity associated with those objects was the same across the two groups. These results are in line with the notion that there are multiple kinds of evaluative information associated with any given stimulus (e.g., Abelson, 1976, 1981; Barsalou, 1992; Bower, 1981; Carlston, 1994; Fishbein & Ajzen, 1975; Fiske & Pavelchak, 1986; Schank & Abelson, 1977; Smith, 1992; Smith & Zarate, 1992), rather than a solitary, summary evaluate tag. Furthermore, these types of

evaluation information might become more or less accessible in orthogonal ways. That is, whereas some types of evaluative information associated with a given object might become more accessible, other types of evaluative information might become more inhibited. The single-tag perspective seems ill-equipped to handle such simultaneous processes (see Chaiken & Bargh, 1993 for a similar argument concerning ambivalent attitudes).

More broadly, the single-tag perspective would also have to argue that every different instantiation of an attitude object is associated with a unique evaluative tag. This would seem to pose a problem of economy as there would have to be an evaluative tag associated with every possible feature, subtype, category, context, and exemplar, not to mention limitless combinations of these variables. As Bassili and Brown have recently argued (Bassili & Brown, 2005), the notion that an attitude exists for every possible instantiation of a particular object-centered-context seems unwieldy and too demanding for a cognitive system.

However, the "dual-tag" view, in comparison, assumes that objects are associated with two attitudes (i.e., evaluative tags), one that becomes activated implicitly, and one that becomes activated explicitly (Wilson et al., 2000). In line with much work in social and cognitive psychology (e.g., Bargh, 1984, 1989, 1990; Jastrow, 1906; Schneider & Fisk, 1982; Shiffrin & Dumais, 1981; Shiffrin & Schneider, 1977; Smith & Lerner, 1986; cf. Gollwitzer, 1996, 1999), this second perspective assumes that an integration of recently acquired evaluative information requires conscious effort. Because of this, implicit measures, which operate outside of the respondent's awareness, presumably can only capture preexisting, stable attitudes (e.g., Wilson & Hodges, 1992; cf. Bargh, 1999; Duckworth et al., 2002). Explicit measures, on the other hand, are assumed to be able to assess more complex integrations of object-relevant and novel information.

Because of the complex integrative processes that can occur during explicit attitude measurement, these researchers argued that the explicit attitude toward a given object is frequently host to a number of contextual parameters such as social desirability, mood, and temporary goals and thus can vary considerably across measurement times. The authors argued that the implicit attitude for any object, however, is quite stable and is not generally amenable to modification. Although one may develop a new explicit attitude toward an object, the old, well-learned, and highly stable implicit attitude continues to exist and may be expressed under some circumstances. For instance, Wilson et al. (2000) asserted that, "Explicit attitudes change relatively more easily whereas implicit attitudes, like old habits, change more slowly. Attitude change techniques often change explicit but not implicit attitudes" (p. 4). It should be noted that, according to this perspective, data from explicit versus implicit *measures* reflect explicit versus implicit *attitudes*.

The dual-tag perspective does not seem consistent with the accumulating evidence of the contextual dependence of automatic attitudes. One could boot-strap this perspective by assuming, as with the single-tag perspective, that contextual dependence arises out of varying categorization, but this runs into the same problem of simultaneous inhibition and excitatory processes that the single-tag view faces. Further, it is not clear why there would only be a single implicit and explicit attitude for every object – why not multiple implicit (or explicit) attitudes? This possibility would seem to suggest that a given attitude object might hypothetically be associated with numerous implicit and explicit attitudes, which would seem an unwieldy solution to the evidence of context-sensitivity.[2]

Both the single-tag and dual-tag perspectives assume that an automatic attitude reflects the evaluation toward the respective object. For the single-tag perspective, this assumption is

consistent with a symbolic model of memory, which assumes that constructs can be retrieved from memory in isolation of other information. Thus, for example, it is possible for an implicit attitude measure to capture the nature of the summary evaluative tag associated with a given object. For the dual-tag perspective, however, implicit attitudes are assumed to emerge out of connectionist, non-symbolic processes (Wilson et al., 2002; see also Gawronski et al., in press; Smith & DeCoster, 1999). Although such a model of memory would necessarily mean that any attitude would reflect a combination of sources of evaluative information, as we describe later, the dual-tag perspective seems to assume otherwise. In fact, the assumption of this perspective is that an implicit attitude is slow to develop and change, and thus should exhibit stability and resistance across contexts. Wilson and colleagues (Wilson et al., 2000, p. 4) assume that implicit attitudes should display much less contextual dependence than explicit attitudes.

In sum, both the single-tag and dual-tag perspectives argue that implicit measures capture attitudes that must be preexisting. In addition, such evaluations are entirely or primarily in reference to the objects in isolation of other kinds of information. These perspectives assume, in other words, that what is measured during an implicit attitude measure is an approximate 1-to-1 mapping of the observed response to the stored attitude toward that object in memory (plus measurement error). According to this perspective, it is possible to measure attitudes stored in memory alongside their associated attitude objects. In other words, it is at least logically possible to directly measure the evaluative tag that is associated in memory with a given attitude object, assuming a lack, or minimum, of measurement error.

The claim that implicitly measured attitudes must be preexisting is in accord with the view in experimental psychology that responses must be repeatedly enacted over time in order to automatically operate, as previously mentioned (Bargh, 1984, 1989, 1990; Fazio, 1986; Jastrow,

1906; Schneider & Fisk, 1982; Shiffrin & Dumais, 1981; Shiffrin & Schneider, 1977; Smith & Lerner, 1986; cf. Bargh, 1999; Duckworth et al., 2002; Gollwitzer, 1996, 1999). This assumption is evident in researcher's attempts to foster new implicitly measured attitudes. For example, in order to change participants' implicitly measured attitudes toward elderly persons, Karpinski and Hilton (2001) exposed participants to over 200 examples of pro-elderly stimuli. The extreme repetition was assumed to be essential for participant's automatic attitudes to shift (see also Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000).

Automatic attitudes toward object-centered-contexts (a constructivist view)

An alternative theoretical possibility is that the observed response on an implicit attitude measure is instead the result of a computation performed by an evaluative system, on the basis of numerous representations, possibly including multiple categories and exemplars that relate to the attitude object in various ways (Bassili & Brown, 2005; Duckworth et al., 2002; E. R. Smith, 1997, 2000; Ferguson, in press; Ferguson & Bargh, 2003). In this way, automatic attitudes would be sensitive to whatever evaluative information is contextually relevant at the time of measurement. Accordingly, the attitude that is implicitly measured does not represent the activation of a single evaluative tag as a function of only the attitude object (Fiske & Pavelchak, 1986) but is instead an integration across multiple sources of affective information.

An integrative process would technically be compatible with the traditional "file-drawer" models of memory in which concepts are represented by nodes, and the nodes are interconnected according to semantic similarity (e.g., see Smith, 1996). In these types of models, each representation is inert and static when it is not activated. These static representations are periodically manipulated by processes such as encoding, storage, and retrieval, for example. Such a model suggests that a single representation can be activated independently of other

activation. It would theoretically be possible for an integrative process to take place across numerous discrete representations.

However, the notion of an integrative process across various sources of evaluative information is directly predicted by parallel distributed processing models of memory (e.g., Anderson & Rosenfeld, 1988; Bechtel & Abrahamsen, 1991; Carlston and Smith, 1996; Smith, 1996; Smith & DeCoster, 1999). Models of such connectionist systems include the assertion that every observable behavior (explicit or implicit) is the result of a fleeting *state* of the mind, wherein all representations are potentially implicated or contributive (see Smith, 1996). An integration across multiple sources of evaluative information would be highly compatible with the assumptions of connectionist systems (see also Arieli, Sterkin, Grinvald, & Aertsen, 1996; Fiedler, 1996).

If attitude generation is supported by connectionist systems, then any observable response would be the result of an integration of information about both the particular object as well as about the context in which the object is encountered. Recent publications have suggested this type of approach for explaining findings of contextual dependence of automatic attitudes (e.g., Bassili & Brown, 2005; Ferguson, 2003; Ferguson & Bargh, 2004; Mitchell et al., 2003) and these have also advocated for connectionist modeling more generally, rather than attempting to use under-specified models of associative networks that are difficult to formalize (e.g., Smith, 1996, 1997).

Object-centered-contexts

According to the evidence of the contextual dependence of automatic attitudes, and the recent theoretical approaches to attitude generation, we argue that a given attitude object can implicitly evoke multiple potential attitudes, according to the context in which the attitude object

is perceived. Implicitly measured attitudes are sensitive to recently acquired information, including goal-relevant expectations, and are each based on an integrative process whereby numerous sources of evaluative information contribute to the observed attitudinal response. To this end, an implicitly measured attitude could be defined as an indication of the relevance of the object for the perceiver as a function of past, current, and future object-relevant circumstances.

Because of the fact that an object is therefore never evaluated in a vacuum, it seems misleading to refer to an implicitly measured attitude as an evaluation of the corresponding *object*. Rather, on the basis of the current argument and results, an implicitly measured attitude corresponds to an *object-centered context*. This new definition is in contrast with how attitudes and attitude objects are typically discussed in attitude literature. For example, Eagly and Chaiken (1993; pp. 4-5) stated that

> "An evaluation is always made with respect to some entity or thing that is the object of the evaluation...Virtually anything that is discriminable can be evaluated and therefore function as an attitude object. Some attitude objects are abstract (e.g., liberalism, secular humanism), and others are concrete (e.g., a chair, a shoe). Particular entities (e.g., my green pen) can function as attitude objects, as can classes of entities (e.g., ballpoint pens). Behaviors (e.g., playing volleyball) and classes of behaviors (e.g., participating in athletic activities) can also function as attitude objects. In general, anything that is discriminated or that becomes in some sense an object of thought can serve as an attitude object."

Based on this excerpt from the widely influential attitudes text by Eagly and Chaiken (1993), it seems possible to measure a solitary attitude toward a solitary object. For example, the excerpt suggests that it is possible to measure one's attitude toward a chair. According to the new definition of an attitude as an indication of an object-based context, however, the best one could do to calculate one's attitude toward chairs *in general* would be to average that person's attitudes toward multiple object-based contexts (a lounge chair, a desk chair, a wooden chair, a chair when someone wants to sit down versus not, etc.). However, the utility of such an average might be suspect because the average attitude across time points might not serve as a very

accurate predictor of behavior in a given particular context. As has long been argued by attitude theorists (e.g., Fishbein & Ajzen, 1974, 1975; Kelman, 1974; Schuman & Johnson, 1976), increasing the specificity with which one describes an attitude object and object-relevant behavior increases the chance that the (explicit) attitude predicts the relevant behavior. The present research suggests that the same reasoning can be applied to implicitly measured attitudes. Thus, it seems potentially useful to conceptualize the referents of automatic attitudes with greater specificity.

Issues related to object-centered-contexts

The concept of attitude stability is a traditionally important issue in attitude research (e.g., Eagly & Chaiken, 1993). There are two points of clarification with regard to stability that are related to the present discussion concerning the contextual dependence of implicitly measured attitudes. The first point of clarification is the distinction between the stability of an attitude, as measured (i.e., observational stability) versus the stability of the representations in memory that underlie the attitude expression (i.e., representational stability). The second point of clarification is the distinction between the concept of contextual dependence and each of the two meanings of stability.

According to a perspective that assumes a 1-to-1 mapping of the observed attitude to the stored attitude in memory, the two meanings of attitude stability (representational versus observational) would be isomorphic except for the issue of measurement noise. That is, given the assumption of 1-to-1 mapping, if researchers observed that an attitude was stable across time points, they might assume that the attitude construct remained unchanged in memory. If one observed attitude instability across time, one could assume that either the attitude construct had changed across time, or that some measurement noise of the instrument was obscuring stability,

or some combination of the two possibilities. Thus, if one assumes the 1-to-1 mapping assumption, then observational and representational stability would be highly related, except for the variance from measurement error.

However, according to the claim that an observed attitude (either explicitly or implicitly measured) reflects an integration of multiple sources of evaluative information, these two meanings of attitude stability become mutually exclusive. That is, because the observed attitude reflects a computation across numerous components, it is impossible to ascertain the individual status of one component of that function, such as the status of the evaluative information associated with one relevant exemplar of the attitude object, for instance. More specifically, even when there is observational stability across time points, it is impossible to infer that the same particular components contributed to the integration, or even in the same way, across time points. It may be that the integration process across time generates a consistently positive attitude, for instance, but this does not indicate anything concerning the representational status of evaluative information associated with each of the components of the integration.

When there is observational instability across time, assuming little measurement noise, it is possible that either the set of particular evaluative memories contributing to the integration has changed, or that the set of evaluative components is the same, but the components themselves have undergone some change. For example, if someone displays a positive attitude toward the elderly on one occasion and a negative attitude toward the same group at a later time, perhaps after some new information, there could be two broad reasons for the instability from Time 1 to Time 2 (not including measurement error). One reason might be that the components involved in the integration changed from Time 1 to Time 2. Assuming a hypothetical array of "elderly" memories that might contribute to the overall integration, whereas elderly memories 4-8

contributed to the integration at Time 1, elderly memories 6-15 might have contributed to the integration at Time 2. A second reason might be that the same elderly memories (4-8) were involved in the integration at both time points, but that the valence of the memories changed. (It should be noted that this possibility assumes that the evaluative information associated with a particular memory can be replaced by different evaluative information.)

The adoption of either the direct mapping versus integration claim influences how freely researchers can generate conclusions about the contents of memory, based on data from attitude measures. Given the assumption of the 1-to-1 mapping, researchers could use attitude measures to generate inferences about the specific organization and contents of memory. For example, if someone demonstrated a positive, implicitly measured attitude toward a certain group, researchers might assume that the perceiver possesses a representation of that group that is associated with a positive evaluative tag. Given the integration claim however, it is no longer viable to make such inferences. It is impossible to postulate about individual pieces of information in memory solely on the basis of an integrative function. Although an object may be associated with a valence in memory and this association may be stable across time, that individual association or valence will always be one component that contributes to the overall summary response that is observed. Thus, according to the integrative perspective, although there may be stability in the memorial representations involving object-related memories and their respective valences, the observable response of a person toward an object will reflect a complex computation, to which many object relevant and irrelevant memories and their associated valences contribute.

Contextual dependence and observational stability

The next consideration is the observational stability of implicitly gauged attitudes. Although the degree of observational stability of an attitude is sometimes interpreted as a proxy for the degree of contextual dependence of the attitude, or vice versa, the current argument is that observational stability and contextual dependence are orthogonal concepts. It is possible that an implicit attitude, even though dependent on the context in which it is measured, can be (observationally) stable or unstable across time.

When will a contextually dependent, implicitly measured attitude be observationally stable across contexts? An object will evoke the same attitude across time points (i.e., demonstrating observational stability) if all object-relevant memories are univalent, if the object-relevant memories that are activated across time points are consistently of the same valence, or if the usefulness of the object (harmful versus helpful) is constant across time. In this way, the stable attitude can be thought of as an indication of stable circumstances or (object-relevant) situational constancy. In contrast, a contextually dependent, implicitly measured attitude will be observationally unstable across contexts if the set of object-relevant memories that is activated across contexts differs in terms of valence, or because the goal-relevance of the object changes across contexts (specifically, from harmful to safe, or vice-versa, for instance).

Observational instability should therefore be somewhat positively related to complex and ambiguous attitude objects. The more object-relevant memories exist, the more likely the object will be able to be perceived in different ways across contexts. As long as some of those object-relevant memories are mixed in valence, the more likely the implicitly invoked attitude will change in direction across time and situations. In general, the more mixed (i.e., both positive and negative) the object-relevant memories, the more likely the implicitly measured attitude will be observationally unstable across time.

Contextual dependence and representational stability

As was the case with observational stability, the contextual dependence of an attitude does not indicate anything about the representational stability underlying that attitude. Representational stability is somewhat difficult to consider as the nature (e.g., organization and format) of representations in memory is still far from clear (e.g., Pinker, 1997; Plotkin, 1997). However, it is possible to speculate that the object and situation-relevant memories that typically comprise an integration can either remain stable or undergo change. (This is assuming that the evaluative information associated with a memory can change, rather than the idea that a change in evaluation forces a whole novel object representation with the new attitude.)

Accordingly, the contextually dependent, implicitly measured attitude can result from integrations that are based on either stable (static) pieces of evaluative information (although not always the same pieces, and not always weighed in the same way), or pieces of evaluative information that undergo change. In short, the contextual dependence of an attitude does not imply anything about its underlying representational stability.

The context-dependence of automatic versus deliberate attitudes

Similar to the argument proposed in the current chapter, researchers have claimed that explicit attitudes are heavily (or entirely) dependent on the context in which the perceiver encounters the corresponding stimuli (e.g., Anderson, 1974; Schwarz & Bohner, 2001; Tesser, 1978). For example, Schwarz and Bohner (2001) identify five steps at which contextual factors might influence the reporting of an attitude, from interpreting the nature of the attitude query to deciding whether to report one's consciously accessible attitude. We wholeheartedly agree that explicit attitudes are likely influenced by a wide array of contextual factors. However, we also assert that the contextual dependence of explicit measures might be somewhat obscured by the

metacognitive and introspective thinking that such measures enable (e.g., Cohen & Schooler, 1997). In particular, because a respondent during an explicit measure has to consciously generate and then report her or his attitude, that person's implicit or even explicit theories about the stability of her or his attitudes might encourage a stability of the attitude that would not otherwise emerge. In line with research by Nisbett and Wilson (1977) among others, people's implicit theories about whether they like pizza, for example, might push them toward reporting a more consistently positive (explicit) attitude toward it, even while their automatic attitude might reflect considerably more fluctuation as a function of current goals, expectations, recently encountered information, and the like.

Another possibility is that automatic attitude measures might tap associations and information that are not consciously accessible to the person. This might mean that a greater wealth of information is potentially more influential on the attitude when the person is not consciously thinking about that attitude. This argument is consistent with recent work by Greenwald and colleagues (Greenwald, Banaji, Rudman, Farnham, Nosek, & Mellott, 2002) that suggests that while implicit measures are subject to cognitive consistency pressures (Heider, 1946, 1958), explicit measures do not seem to be. For example, this work suggests that when a woman displays automatic positivity toward herself, and toward females in general, she will likely not display automatic positivity toward math, which is assumed to be dissociated from the concept of females (see also Nosek et al., 2002). Implicitly measured attitudes, from this perspective, may be sensitive to and influenced by a more extensive set of memories than are explicitly measured attitudes. Future research can undoubtedly shed more light on the comparative degree to which implicit versus explicit measures are dependent on the context in

which they are administered, but we put forth at this point the possibility that implicit (versus explicit) measures yield more contextually-sensitive attitudes.

Conclusion

Research and theory conducted and developed over the last 5 years suggest that automatic attitudes are highly variable on numerous contextual factors. Specifically, research findings suggest that automatic attitudes vary depending on current goals, task instructions, recently encountered attitude objects and attitude-relevant contexts, and even mood. In line with this evidence, recent theoretical approaches have argued that an automatic attitude is likely the result of an integrative process within a connectionist system. In such a system, the pattern of activation across a set of processing units for any given attitude is re-created each time the attitude object is encountered. Although there are undoubtedly stable sources of evaluative information reflected in the weights that connect the processing units, the resulting attitude is nevertheless based on a constructive process that incorporates multiple sources of information. In this way, slight changes in the type of evaluative information included in the computation according to the context can potentially change the direction or intensity of the automatic attitude. This research and theory suggest that automatic attitudes should be understood as evaluations of object-centered-contexts, rather than objects in isolation.

Importantly, although implicit measures clearly do not assess attitudes free of the context in which the objects are perceived, they are certainly nevertheless valuable research tools. Implicit attitude measures can effectively gauge a person's response tendency (either approach if the attitude is positive or avoidance if the attitude is negative) within certain situations, with regard to a particular object. This is useful to the extent that researchers can show that such implicitly gauged responses are predictive of the person's behavior (either subtle or overt) in a

future similar situation. A large body of research now has shown that implicit attitude measures can predict a variety of both nonverbal and overt behaviors toward the attitude object (e.g., for reviews see Blair, 2002; Fazio & Olson, 2003; Ferguson, in press; Lane et al., this volume; Wittenbrink, this volume).

Furthermore, implicit attitude measurement is also useful because it can indicate various characteristics of nonconscious processing. Whereas the initial assumption in attitude research and in social psychology in general was that automatic responses can only be the result of repeated previous experiences with the object (e.g., Bargh, 1984, 1989, 1990; Jastrow, 1906; Schneider & Fisk, 1982; Shiffrin & Dumais, 1981; Shiffrin & Schneider, 1977; Smith & Lerner, 1986; cf. Gollwitzer, 1996, 1999), and cannot reflect recent integrations of new information (e.g., Wilson & Hodges, 1992; Wilson et al., 2000), recent research has overturned these assumptions (see Bargh, 2001; Dijksterhuis, 2004; Duckworth et al., 2002). There is still much to learn about people's ability to automatically assess their environments, and implicit attitude measures represent one valuable tool to accomplish this objective.

[1] The terms implicit attitudes and explicit attitudes should be assumed to reflect the type of measurement, rather than the type of attitude (see DeHouwer, in press; Fazio & Olson, 2003).

[2] In addition, the proposal that implicit attitudes are distinct from explicit attitudes is not consistent with research suggesting that they are in fact related constructs (Hofman, Gawronski, Gschwendner, Le, & Schmitt, in press; Nosek, in press).

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.

Albarracín, D., Johnson, B. T., & Zanna, M. P. (2005). *The Handbook of Attitudes*. Mahwah, NJ: Erlbaum.

Anderson, J. R., & Bower, G. H. (1973). *Human associative memory*. Washington, D.C.: Winston and Sons.

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.

Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science, 273,* 1868-1871.

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.

Barden, J., Maddux, W. W., Petty, R.E., & Brewer, M.B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology, 87*, 5-22.

Bargh, J. A. (1984). Automatic and conscious processing of social information. In R. S. Wyer Jr., & T. K. Srull (Eds.), *Handbook of social cognition* (Vol. 3, pp. 1-43). Hillsdale, NJ: Erlbaum.

Bargh, J. A. (1989). Conditional automaticity: Varieties of automatic influence in social perception and cognition. In J. S. Ulerman & 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, NY, USA: The Guilford Press.

Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. J. Wyer & T. K. Srull (Eds.), Handbook of social cognition (pp. 1-40). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer. Jr. (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1-61), Mahwah, NJ: Erlbaum.

Bargh, J. A. (1999). The cognitive monster. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 361-382). New York: Guilford.

Bargh, J. A. (2001). Caution: Automatic social cognition may not be habit forming. *Polish Psychological Bulletin, 32,* 1-8

Bargh, J. A., Chaiken, S., Govender, 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., & Hymes, C. (1996). The automatic evaluation effect: Unconditional automatic attitude activation with a pronunciation task. *Journal of Experimental Social Psychology, 32,* 104-128.

Barsalou, L. W. (1992). *Cognitive psychology: An overview for cognitive scientists.* Hillsdale, NJ: Erlbaum.

Bassili, J. N., & Brown, R. D. (2005). Implicit and explicit attitudes: Research, Challenges, and theory. In D. Albarracín, B. T. Johnson & M. P. Zanna (Eds.), *The Handbook of Attitudes* (pp. 543-574). Mahwah, NJ: Erlbaum.

Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind: An introduction to parallel processing in networks.* Oxford, England UK: Basil Blackwell, Inc.

Bem, D. J. (1972). Constructing cross-situational consistencies in behavior: Some thoughts on Alker's critique of Mischel. *Journal of Personality, 40,* 17-26.

Blair, I. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychological Review, 6*, 242-261.

Blair, I. V., Ma, J. E., & Lenton, A. P. (2001). Imagining stereotypes away: The moderation of implicit stereotypes through mental imagery. *Journal of Personality and Social Psychology, 81*, 828 – 841.

Bower, G. H. (1981). Mood and memory. *American Psychologist, 36,* 129-148.

Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin, 115,* 401-423.

Cacioppo, J. T., Gardner, W. L., & Berntson, G. G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review, 1*, 3-25.

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, NY: Guilford Press.

Carlston, D. E. (1994). Associated systems theory: A systematic approach to the cognitive representation of persons and events. In R. S. Wyer (Ed.) *Advances in Social Cognition: Vol. 7. Associated Systems Theory* (pp. 1-78). Hillsdale, NY: Erlbaum.

Chaiken, S., Bargh, J. A. (1993). Occurrence versus moderation of the automatic attitude activation effect: Reply to Fazio. *Journal of Personality & Social Psychology, 64*, 759-765.

Chaiken, S., & Yates, S. (1985). Affective-cognitive consistency and thought-induced attitude polarization. Journal of Personality and Social Psychology, 49, 1470-1481.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review, 82,* 407-428.

Cohen, J. D., & Schooler, J. W. (1997). Scientific approaches to consciousness. Hillsdale, NJ: Lawrence Erlbaum Associates.

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.

De Houwer, J., & Eelen, P. (1998). An affective variant of the Simon paradigm. *Cognition and Emotion, 8*, 45-61.

De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology, 50*, 77-85.

De Houwer, J. (in press). What are implicit measures and why are we using them? In R. W. Wiers & A. W. Stacy (Eds.), *The handbook of implicit cognition and addiction*. Thousand Oaks, CA: Sage Publishers.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56,* 5-18.

Dijksterhuis, A. (2004). I like myself but I don't know why: Enhancing implicit self-esteem by subliminal evaluative conditioning. *Journal of Personality and Social Psychology, 86*, 345-355.

Dovidio, J. F., Kawakami, K., Johnson, C., & Johnson, B. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology, 33,* 510-540.

Duckworth, K. L., Bargh, J. A., Garcia, M., & Chaiken, S. (2002). The automatic evaluation of novel stimuli. *Psychological Science, 13*, 513 - 519.

Eagly, A. H., & Chaiken, S. (1993). *The psychology of attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich College.

Fazio, R. H. (1987). Self-perception theory: A current perspective. In M. P. Zanna, J. M. Olson, & P. C. Herman (Eds.*) Social influence: The Ontario symposium, Vol. 5* (pp. 129-150). Hillsdale, NJ: Erlbaum.

Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty, & J. A. Krosnick (Eds.) *Attitude strength: Antecedents and consequences* (pp. 247-282). Hillsdale, NJ: Erlbaum.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297-327.

Fazio, R. H., & Williams, C. J. (1986). Attitude accessibility as a moderator of the attitude-perception and attitude-behavior relations: An investigation of the 1984 presidential election. *Journal of Personality and Social Psychology, 51*, 505-514.

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., & Williams, C. J. (1986). Attitude accessibility as a moderator of the attitude-perception and attitude-behavior relations: An investigation of the 1984 presidential election. *Journal of Personality and Social Psychology, 51*, 505-514.

Fazio, R. H. (1986). How do attitudes guide behavior? In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 204-243). New York: Guilford.

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.153-179). Hillsdale, N.J.: Erlbaum.

Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 75-109). New York: Academic Press.

Fazio, R. H. (2001). On the automatic activation of associated evaluations: An overview. *Cognition and Emotion, 14,* 1-27.

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.

Feldman, J. M., & Lynch, J. G. (1988). Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of Applied Psychology, 73*, 421-435.

Ferguson, M. J. (in press). The automaticity of evaluation. Invited chapter to appear in J. A. Bargh (Ed.), *Automatic Processes in Social Thinking and Behavior*. Psychology Press.

Ferguson, M. J. (2005). The automatic evaluation of end-states: How automatic attitudes toward goals predict and influence judgment and behavior. Unpublished manuscript. Cornell University.

Ferguson, M. J., & Bargh, J. A. (2003). The constructive nature of automatic evaluation. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp.169-188). NJ: Erlbaum.

Ferguson, M. J., & Bargh, J. A. (2004). Liking is for doing: Effects of goal pursuit on automatic evaluation. *Journal of Personality and Social Psychology, 88*, 557-572.

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 towards 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., & 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.

Fiske, S. T. (1992). Thinking is for doing: Portraits of social cognition from Daguerreotype to laserphoto. Journal of Personality and Social Psychology, 63, 877-889.

Forgas, J. P. (1992). Affect in social judgments and decisions: A multiprocess model. In M. P. Zanna (Ed.) *Advances in experimental social psychology, Vol. 25* (pp. 227-275). San Diego, CA: Academic Press.

Gawronski, B., Strack, F., & Bodenhausen, G. V. (in press). Attitudes and cognitive consistency: The role of associative and propositional processes. To appear in R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new wave of implicit measures*.

Gollwitzer, P. M. (1996). The volitional benefits of planning. In P. M. Gollwitzer & J. A. Bargh (Eds.), *The psychology of action: Linking cognition and motivation to behavior* (pp. 287-312). New York, NY: Guilford Press.

Gollwitzer, P. M. (1999). Implementation intentions: Strong effects of simple plans. *American Psychologist, 54,* 493-503.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102,* 4-27.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review, 109*, 3-25.

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., McGhee, 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.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology: Interdisciplinary and Applied, 21,*107-112.

Heider, F. (1958). *The psychology of interpersonal relations*. New York: Wiley

Hermsen, B., Holland, R. W., & van Knippenberg, A. (2004). The influence of mood on automatic evaluations. Unpublished Manuscript. University of Nijmegen.

Hofmann, W., Gawronski, B., Gschwendner, T., Le, H., & Schmitt, M. (in press). A meta-analysis on the correlation between the Implicit Association Test and explicit self-report measures. *Personality and Social Psychology Bulletin*.

James, W. (1890). The principles of psychology (Vol. 2). New York: Holt.

Jastrow, J. (1906). *The subconscious*. Boston, MA: Houghton-Mifflin.

Jones, E. E., & Sigall, H. (1971). The bogus pipeline: A new paradigm for measuring affect and attitude. *Psychological Bulletin, 76,* 349-364.

Karpinski, A., & Hilton, J. L. (2001). Attitudes and the implicit association test. *Journal of Personality and Social Psychology, 81*, 774-788.

Kawakami, K., Dovidio, J.F., Moll, J., Hermsen, S., & Russin, A. (2000). Just say no (to stereotyping): Effects of training in the negation of stereotypic associations on stereotype activation. *Journal of Personality and Social Psychology, 78*, 871-888.

Kelman, H. C. (1974). Attitudes are alive and well and gainfully employed in the sphere of action. *American Psychologist, 29,* 310-324.

Koole, S.K., Dijksterhuis, A., & van Knippenberg, A. (2001). What's in a name: Implicit self-esteem. *Journal of Personality and Social Psychology, 80*, 614-627.

LaPiere, R. T. (1934). Attitudes vs. actions. *Social Forces, 13,* 230-237.

Livingston, R. W. & Brewer, M.B. (2002). What are we really priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology, 82*, 5-18.

Lowery, B.S., Hardin, C. D., & Sinclair, S. (2001). Social influence on automatic racial prejudice. *Journal of Personality and Social Psychology, 81*, 842-855.

Maddux, W. W., Barden, J., Brewer, M. B., & Petty, R. E. (2005). Saying no to negativity: The effects of context and motivation to control prejudice on automatic evaluative responses. *Journal of Experimental Social Psychology, 41*, 19-35.

McCauley, C., & Stitt, C. L. (1978). An individual and quantitative measure of stereotypes. *Journal of Personality and Social Psychology, 36,* 929-940.

McConahay, J. B. (1983). Modern racism and modern discrimination: The effects of race, racial attitudes, and context on simulated hiring decisions. *Personality and Social Psychology Bulletin, 9,* 551-558.

McConahay, J. (1986). Modern racism, ambivalence, and the Modern Racism scale. In J. Dovidio (Ed.), *Prejudice, discrimination, and racism* (pp. 91-125). San Diego, CA: Academic Press.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90,* 227-234.

Miller, M. G., & Tesser, A. (1986). Thought-induced attitude change: The effects of schema structure and commitment. *Journal o Personality and Social Psychology, 51,* 259–269

Mitchell, J. P., Nosek, B. A., & Banaji, M. R. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General, 132*, 455-469.

Moors, A., & De Houwer, J. (2001). Automatic Appraisal of Motivational Valence: Motivational Affective Priming and Simon Effects. *Cognition and Emotion, 15*, 749-766.

Moors, A., De Houwer, J., & Eelen, P. (2004). Automatic stimulus-goal comparisons: Evidence from motivational affective priming studies. *Cognition and Emotion, 18*, 29-54.

Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology, 77*, 167-184.

Musch, J., & Klauer, K. C. (2003). *The Psychology of Evaluation: Affective Processes in Cognition and Emotion*. Mahwah, NJ: Lawrence Erlbaum.

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for faciliatory 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.

Nisbett, R., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*, 231-259.

Nosek, B. A., & Banaji, M. R. (2001). The go/no-go association task. *Social Cognition, 19(6)*, 625-666.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002). Math=Me, Me=Female, therefore Math is not equal to me. *Journal of Personality and Socib*xternal motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811-832.

Nosek, B. A. (in press). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*.

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. *American Psychologist, 17*, 776-783.

Payne, B.K., Cheng, C.M., Govorun, O., & Stewart, B. (in press). An inkblot for attitudes: Attitude misattribution as implicit measurement. *Journal of Personality and Social Psychology*.

Petty, R. E., & Krosnick, J. A. (1995). *Attitude strength: Antecedents and consequences*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Pinker, S. (1997). *How the mind works*. New York, NY: W. W. Norton and Co, Inc.

Plotkin, H. 1997. *Evolution in mind*. London: Penguin.

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition: The Loyola symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.

Rudman, L. A., Ashmore, R. D., & Gary, M. L. (2001). "Unlearning" automatic biases: The malleability of implicit stereotypes and prejudice. *Journal of Personality and Social Psychology, 81*, 856-868.

Salancik, G. R., & Conway, M. (1975). Attitude inferences from salient and relevant cognitive content about behavior. *Journal of Personality and Social Psychology, 32,* 829-840.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Erlbaum.

Schneider, W., & Fisk, A. D. (1982). Degree of consistent training: Improvements in search performance and automatic process development. *Perception & Psychophysics, 31,* 160-168.

Schuman, H., & Johnson, M. P. (1976). Attitudes and behavior. *Annual Review of Sociology, 2,* 161-207.

Schwarz, N., & Bless, H. (1992). Constructing reality and its alternatives: Assimilation and contrast effects in social judgment. In L.L. Martin & A. Tesser (Eds.), *The construction of social judgment* (pp. 217-245). Hillsdale, NJ: Erlbaum.

Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Blackwell handbook of social psychology: Intraindividual processes* (Vol.1, pp. 436-457). Oxford, UK: Blackwell.

Schwarz, N., & Clore, G. L. 1983. Mood, misattribution, and judgments of well-being, Informative and directive functions of affective states. *Journal of Personality and Social Psychology, 45,* 513-523.

Sherman, S.J.**,** Presson, C.C., Chassin, L., Rose, J.S., & Koch, K. (2003). Implicit and explicit attitudes toward cigarette smoking: The effects of context and motivation. *Journal of Social and Clinical Psychology, 22*, 13-39.

Shiffrin, R. M., & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp.111-140). Hillsdale, NJ: Erlbaum.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review, 84,* 127-190.

Schuman, H., Presser, S. (1981). *Questions and Answers in Attitude Surveys: Experiments on Question Form, Wording, and Context*. New York: Academic.

Smith, E. R. (1992). The role of exemplars in social judgment. In L. L. Martin and A. Tesser (Eds.), *The construction of social judgment*. Hillsdale, NJ: Erlbaum.

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. (1997). Preconscious automaticity in a modular connectionist system. In R. S. Wyer (Ed.), *Advances in social cognition* (Vol. 10, pp. 181-202). Mahway, NJ: Erlbaum.

Smith, E. R. (2000, February). *Connectionist representation of evaluation*. Paper presented at the meeting of the Society for Personality and Social Psychology, Nashville, TN.

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, NY: Guilford Press.

Smith, E. R., & Lerner, M. (1986). Development of automatism of social judgments. *Journal of Personality and Social Psychology, 50,* 246-259.

Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99,* 3-21.

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. E., Rudman, L. A., & Greenwald, A. G. (2001). Using the Implicit Association Test to investigate attitude-behaviour consistency for stigmatised behaviour. *Cognition and Emotion, 15,* 207-230.

Rosenberg, M. J. (1969). The conditions and consequences of evaluation apprehension. In R. Rosenthal & R. Rosnow (Eds.), *Artifact in behavioral research* (pp. 279-349). NY: Academic Press.

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.

Tourangeau, R. & Rasinski, K.A. (1988). Cognitive processes underlying context effects in attitude measurement. *Psychological Bulletin, 103*, 299-314.

von Hippel, W., Sekaquaptewa, D., & Vargas, P. (1997). The Linguistic Intergroup Bias as an implicit indicator of prejudice. *Journal of Experimental Social Psychology, 33,* 490-509.

Weber, R., & Crocker, J. (1983). Cognitive processes in the revision of stereotypic beliefs. *Journal of Personality and Social Psychology, 45,* 961–77.

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.