Published in final edited form as:
Soc Cogn. 2007 October 1; 25(5): 603–637. doi:10.1521/soco.2007.25.5.603.

Attitudes as Object-Evaluation Associations of Varying Strength

Russell H. Fazio
Ohio State University

Abstract

Historical developments regarding the attitude concept are reviewed, and set the stage for consideration of a theoretical perspective that views attitude, not as a hypothetical construct, but as evaluative knowledge. A model of attitudes as object-evaluation associations of varying strength is summarized, along with research supporting the model's contention that at least some attitudes are represented in memory and activated automatically upon the individual's encountering the attitude object. The implications of the theoretical perspective for a number of recent discussions related to the attitude concept are elaborated. Among these issues are the notion of attitudes as "constructions," the presumed malleability of automatically-activated attitudes, correspondence between implicit and explicit measures of attitude, and postulated dual or multiple attitudes.

For nearly 25 years now, a particular view of attitudes has formed the underpinnings of my research program on the consequences of attitudes for attention, categorization, judgment, and behavior. It was in 1982 that my colleagues and I first proposed that attitudes can be viewed as object-evaluation associations in memory (Fazio, Chen, McDonel, & Sherman, 1982). The perspective has proven much more illuminating (and occupied us for much longer) than we had envisioned at the time. It has fostered examination of a wide variety of questions regarding attitudes, and it continues to provide a valuable perspective for considering new issues. In this article, I summarize the theoretical model and some of the research findings that support it, as well as the perspective the model provides regarding some recent discussions of attitudes as constructions, the presumed malleability of attitudes, correspondence between implicit and explicit measures of attitude, and postulated dual or multiple attitudes.

Background

A Very Brief History

Even as long ago as 1935, when Gordon Allport wrote his influential chapter for the *Handbook of Social Psychology*, the term "attitude" already had a rich history. Allport provides an intriguing historical account of the term's early meanings. He notes that "attitude" had been used in the arts to refer to the posture of a figure in a painting or sculpture. This connotation was evident in what was among the very first usages of the term in experimental psychology – within the study of reaction time. In the late 1800's, numerous studies revealed that participants who were mentally prepared to press a telegraph key upon a signal were able to respond more quickly than those whose attention was focused on the incoming signal itself. The importance of this state of preparedness, sometimes referred to as a "task-attitude" or "mental set," was demonstrated repeatedly in studies of perception and memory.

Beginning with these early tracings of the attitude concept, Allport (1935) reviewed a large number of definitions that had been offered. His analysis of their common threads and the

Address correspondence to: Russell H. Fazio, Department of Psychology, 1835 Neil Avenue, Ohio State University, Columbus, Ohio 43210-1287, Email: E-mail: Fazio.11@osu.edu, Phone: 614-688-5408, Fax: 614-688-5414.

debates they had inspired led Allport to propose what is certainly the most widely-known of the early definitions of attitude. "An attitude is a mental and neural state of readiness, organized through experience, exerting a directive or dynamic influence upon the individual's response to all objects and situations with which it is related" (p. 810). Thus, the early references to posture and readiness were featured in Allport's definition. But, the definition seems to go much further when it posits that this state of readiness influences a broad array of responses, which raises an important question. If no such influence is observable, then is there no attitude?

Therein lies what later came to prove especially controversial about the attitude concept. Allport's definition, like many others in the literature, presupposes behavioral responses that are consistent with the attitude. That is, the very definition calls for behavioral consistency. Without evidence of such consistency, the defining criteria of an attitude have not been achieved. In the late 1960s and early 1970s, this assumption regarding behavioral correspondence began to be examined – and seriously challenged. Wicker's (1969) influential review of the available evidence suggested that verbal reports of attitudes bore little relation to subsequent behavior. Bem (1972) espoused a view of attitudes, not as potent determinants of responses, but as epiphenomenal explanations for past behavior. Many attitude researchers rose to the challenge and demonstrated that sizeable relations could be observed between verbal reports of attitude and subsequent behavior under certain conditions. This is not the time to review those findings. Many pertinent reviews are available (e.g., Fazio & Roskos-Ewoldsen, 2005; Wallace, Paulson, Lord, & Bond, 2005; Zanna & Fazio, 1982), all of which clearly confirm that attitudinal reports can sometimes prove strongly predictive of behavior.

More relevant to the present concerns is what this controversy regarding the attitude-behavior relation implies for a definition of attitude. The fact that the answer to the question "Is there a relation between attitudes and behavior?" is a resounding "sometimes" (see Zanna & Fazio, 1982, for an overview of generations of research concerning attitude-behavior consistency) suggests that assumptions regarding behavioral correspondence should not be a part of one's very definition of attitude. A definition that includes such a criterion is open to circularity; it essentially "defines away" the question of attitude-behavior consistency. If behavioral correspondence is not observed, there is no attitude. Hence, the very definition does not allow for the possibility that a multitude of factors may play a role in determining whether an individual's attitude is evident behaviorally in a given situation. Instead of prejudging the attitude-behavior relation, any definition of attitude should leave the matter open to theoretical and empirical investigation (Zanna & Rempel, 1988). This reasoning was very central to the model of attitudes that I came to pursue and will explicate shortly.

Attitudes: Hypothetical Constructs?

Yet there is an additional sense in which a definition of attitude that postulates influences on observable responses at its very core is problematic. Is an "attitude" being viewed as a hypothetical construct? And, if so, "hypothetical" in what sense: hypothetical in the sense that it is not directly observable, or in the sense of being merely a convenient conceptual abstraction for describing observed covariance? Does an attitude "exist" as a latent variable that is presumed to play a mechanistic role when the individual encounters or considers the attitude object? Or, is it no more than a useful codeword for scientific discourse?

Classic discussions of the distinctions that might be drawn between various forms of hypothetical constructs and intervening variables include MacCorquodale and Meehl (1948) and Hilgard (1958). However, the distinction I wish to emphasize here is most clearly articulated by DeFleur and Westie (1963) in their influential article entitled "Attitude as a Scientific Concept." They contrast two different "types of latent process conceptions [of attitude]: (1) those which impute the empirical existence of a hidden mechanism, and (2) those which postulate a hypothetical mediating variable which is not regarded as having empirical

referents, counterparts, or existence but which is simply a construction which serves as a convenient tool for analysis" (p. 24).

To exemplify the latter category, DeFleur and Westie point to Green's (1954) conception:

> Like many psychological variables, attitude is a hypothetical or latent variable, rather than an immediately observable variable. The concept of attitude does not refer to any one specific act or response of an individual, but is an abstraction from a large number of related acts or responses. For example, when we state that individual A has less a favorable attitude toward labor organizations than individual B, we mean that A's many different statements and actions concerning labor organizations are consistently less favorable to labor than are B's comparable words and deeds. We are justified in using a comprehensive concept like attitude when the man's related responses are consistent (p. 335).

Conceptualized in this way, then, "attitudes" do not exist as entities within the individual. Moreover, scientific reference to "attitude" is not even appropriate without evidence of consistent verbal and behavioral responding. Thus, this approach once again raises the issue of prejudging the attitude-behavior relation and "defining away" the problem of inconsistency. To the extent that inconsistency is observed between verbal reports and behavior, there is no "attitude" (or at least should not be), not even in the mind of the scientist.

In contrast, DeFleur and Westie (1963) classify Allport's definition in the first of the two categories they delineated. Because Allport refers to attitudes as "exerting a directive influence," they characterize his definition as including "the *additional* idea that the individual's behavior is somehow 'shaped,' 'guided,' or 'mediated' by some *underlying process* [emphasis in original]…some 'inner mechanism,' some unobservable 'something' that constrains, influences, mediates, or otherwise determines that consistency will appear among the individual's responses to the attitude stimulus" (p. 23). As the tone implies, DeFleur and Westie are critical of the approach, emphasizing two reasons for their skepticism. The first concerns conceptual imprecision and the lack of relevant empirical evidence: "*…the latent process conception of attitude must be entertained as most tentative* [emphasis in original]. In posing such an entity it is not precisely clear, at least at the present time, whether the advocates have named a discovery, or discovered a name. Until the mechanisms involved in this internal latent process can be described more fully, the latent variable must remain 'an unknown something'" (p.24). The second point of skepticism concerns what DeFleur and Westie called "the fallacy of expected correspondence," i.e., the presumption that both verbal reports and actions would be affected by the same latent process. They questioned this presumption on the basis of emerging evidence of inconsistency between verbalizations and other forms of action toward the attitude object (e.g., LaPiere, 1934).

Although the precise meaning that is intended is sometimes not clear, references to attitudes as hypothetical constructs continue to appear in the contemporary literature. For example, Greenwald and Nosek (in press) assert: "Attitudes, like other psychological constructs, are hypothetical and unobservable." Schwarz & Bohner (2001) are more explicit when they characterize attitudes as a scientific reification: "Attitudes are a hypothetical construct, invented by researchers to account for a body of phenomena" (p. 438). In my view, and as I hope will become all the clearer shortly, this lingering ambiguity about the sense in which the attitude concept might be viewed as a hypothetical construct has contributed to the very confusion that inspired this special issue in the first place. To preview the argument briefly, a view of attitudes as hypothetical abstractions that the scientific observer imputes upon targets of observation invites neglect of critical questions regarding the acquisition and use of evaluative knowledge. A functional and adaptive system for judgment and decision-making requires that individuals learn from prior experience, that the acquired evaluative knowledge

be represented in memory, and that it be activated efficiently when it may prove useful in later situations. The science is not positioned to ask these questions if attitudes are viewed as a mere reification. Unless attitudes are thought to have some form of existence within the individual, we fail to acknowledge the importance of prior learning for current evaluative judgment.

Lessons Learned?

Awareness of this background regarding the history of the attitude concept should sensitize any attitude theorist and researcher to certain issues. An adequate conceptualization of attitudes must satisfactorily address the matters that have proven pivotal over the last 75 years or so. To ignore the lessons that have been (should have been) learned detracts from the cumulative advances and progress that are so evident in the literature. First, there appears little value to defining attitudes in such a way as to rest the concept on the fulcrum of behavioral consistency. As Zanna and Rempel (1988) argued, definitions that presuppose attitude-behavior consistency dismiss important questions that should be open to theoretical consideration and empirical investigation. Second, and relatedly, DeFleur and Westie's (1963) "fallacy of expected correspondence" must be avoided. Any theoretical model of attitudes will need to acknowledge that behavior sometimes corresponds to verbal reports of attitude and sometimes does not. Third, unless the idea that attitudes can exert influence is to be abandoned, an adequate conceptualization of attitudes will need to embodied within a model of the process(es) by which such influence occurs. Just as DeFleur and Westie (1963) argued, unobservable "somethings" that remain unspecified hidden mechanisms of influence should not be tolerated. Moreover, the specified mechanisms should contribute to our understanding of why measures of attitude sometimes prove predictive of later behavior and sometimes not.

Attitudes as Object-Evaluation Associations

With this background in mind, let me summarize the view of attitudes that has been the focus of the research my colleagues and I have conducted. The model was first proposed by Fazio et al. (1982), tested in a variety of ways in subsequent research over the next decade, and most systematically delineated in a review chapter by Fazio (1995). In brief, the model views attitudes as associations between a given object and a given summary evaluation of the object – associations that can vary in strength and, hence, in their accessibility from memory.

Thus, attitudes are defined as summary evaluations. However, this should not imply that attitudes are necessarily cold, belief-based judgments of favorability. The term "evaluation" is used broadly to include not only analytic assessments but also "hot" affective reactions. Like Zanna and Rempel's (1988) formulation, the model views these evaluative summaries as potentially stemming from beliefs, affect, and/or behavioral information. The attitude may be based on appraisals of the attributes that characterize the object, as in expectancy-value frameworks. It can stem from emotional reactions that the attitude object evokes, as in the case of conditioned emotional responses. It can be based on one's past behaviors and experiences with the object. Or, it can be based on some combination of these potential sources of evaluative information. Relatedly, the attitude may be the outcome of a rather passive associative learning process (e.g., De Houwer, Thomas, & Baeyens, 2001; Olson & Fazio, 2001). It may be a relatively simple inference from attitudinally-diagnostic information, including freely-chosen behavior (Bem, 1972), proprioceptive feedback (e.g., Strack, Martin, & Stepper, 1988), perceptual fluency (e.g., Reber, Winkielman, & Schwarz, 1998), or ease of retrieval (e.g., Briñol, Petty, & Tormala, 2006). Or, the attitude may represent the summary product of a more active process of propositional reasoning that involves careful scrutiny of the validity of information regarding the attitude object (e.g., Deutsch & Strack, 2006; Gawronski & Bodenhausen, 2006; Petty & Cacioppo, 1986). The processes by which the attitude was formed and its informational basis may have implications for the resulting strength of the object-evaluation association and, hence, the accessibility of the attitude from memory (see Fazio,

1995; Petty, Haugtvedt, & Smith, 1995). However, regardless of how it was formed, the attitude is defined as the summary evaluation, devoid of any presumptions regarding behavioral consistency.

Importantly, the model of attitudes is very explicit with respect to the matter of attitudes being viewed as hypothetical constructs invented by the scientific community. By viewing attitudes as associations in memory, the model obviously is positing that attitudes can "exist." That is, if the individual has formed any such evaluative association, then the attitude is represented in memory. Attitudes then are simply a form of knowledge – evaluative knowledge, more precisely – and are represented in memory just as any other knowledge is represented. Just as we associate "bread" with "butter" and "doctor" with "nurse," we can associate "yuck" with cockroaches, or a sense of ecstatic delight with chocolate or single malt scotch.

Thus, according to this view, attitudes are no more hypothetical than any other form of knowledge. We do not speak of a person's knowledge that ants are small, but strong (i.e, the associations between ant and small and between ant and strong) as a hypothetical construct. Obviously, there is an element of unobservability here, just as is true for attitudes. As scientific observers, we have difficulty discerning whether the individual possesses knowledge regarding the size of ants. We need to pose a verbal question or observe a relevant behavior, such as the individual's muscular preparation when asked to pick up a cardboard box in which an ant has been placed. Likewise, we cannot see the individual's attitude toward single malt scotch directly, but we can inquire, we can observe, and we can arrange tests assessing the individual's liking. In both cases, we can reach the conclusion that the knowledge, evaluative knowledge in the latter case, exists.

Another critical aspect of the model is that it highlights variability in the strength of the object-evaluation association. Some attitudes for some individuals are characterized by strong associations. Others are weaker. And, for still others, no summary evaluation may even be available in memory, possibly because the object falls in a domain that carries no hedonic significance for the particular individual. So, a college basketball fan will have strong evaluative associations to a wide number of teams, coaches, and players, will develop attitudes toward new players readily, and will experience affective reactions to certain teams' wins and losses. On the other hand, for other individuals, evaluative knowledge regarding all but the most heralded entities in the college basketball scene will be absent. In many cases, no a priori evaluation will be available in memory. Some individuals may exhibit a greater propensity to form attitudes across a variety of domains than other individuals do, as is apparent, for example, with respect to the individual difference measure of need to evaluate (Jarvis & Petty, 1996). However, substantial variability across domains is to be expected, just as a function of people's interests. Sports are hedonically relevant for some people but not all. Politics can be rife with evaluative associations for some, but again not for all. Some people have sophisticated evaluative knowledge regarding cooking techniques, recipes, and spices, whereas for others these entities are unlikely to be associated with anything other than indifference and may even be complete unknowns.

Thus, across people and objects, attitudes necessarily vary with respect to their associative strength. Building upon a distinction highlighted by Converse (1970), we have referred to this variability as the *attitude-nonattitude continuum* (Fazio, Sanbonmatsu, Powell, & Kardes, 1986). At the non-attitude end of the continuum is the case of the individual lacking any a priori evaluative association to the object. Due to its novelty or its basis in a sphere of indifference for the individual, no relevant attitudinal representation is available in memory. As we move along the continuum, an evaluation is not only available but also is more strongly associated with the attitude object. Thus, for some people, the object-evaluation association may be such

that the object is attitude-evoking. The associated evaluation is activated automatically from memory upon mere observation or mention of the attitude object.

Given that the term "automatic" has come to be used in multiple senses (see Bargh, 1994; Moors & De Houwer, 2006; Shiffrin & Dumais, 1981, for relevant discussions), it will be useful to spend a moment explicating just what is meant by "automatic" attitude activation within our theoretical model. From the model's perspective, the key feature is inescapability. Encountering the attitude object activates the associated evaluation without the individual's intent, and does so even if the individual is attempting to engage in some other activity. Hence, what is most relevant to our use of the term is Shiffrin and Dumais' (1981) characterization of automatic as "all processes whose initiation the subject cannot control" – ones that occur "whenever a given set of external initiating stimuli are presented, regardless of a subject's attempt to ignore or bypass the distraction" (p. 117). When the object is encountered, the attitude is activated, not necessarily to a level of awareness, although that is possible, but to a level of accessibility that increases the likelihood that the evaluation will influence the interpretation of subsequent information (Bruner, 1957; Higgins, 1996).

Any such automatic activation of the attitude is viewed as playing a critical role in the process by which an attitude may exert influence on information processing, judgment, and behavior. Indeed, from its outset, this theoretical conceptualization of attitudes has been embodied within a model attempting to specify the process(es) by which attitudes "guide" behavior (for early discussions, see Fazio, Powell, Herr, 1983; Fazio, 1986). Thus, the MODE model (Fazio, 1990; Fazio & Towles-Schwen, 1999) postulates that attitudes can influence behavior in a relatively spontaneous fashion, to the extent that they are activated automatically from memory upon the individual's encountering the attitude object. By biasing perceptions of the object in the immediate situation, the activated attitude can prompt attitudinally-consistent behavioral responses. Moreover, this can occur without the individual necessarily engaging in any effortful reflection regarding his or her attitude toward the object and without any necessary awareness on the part of the individual that his or her activated attitude has biased construal of the object.

More will be said about attitude-to-behavior processes in a later section of this article. For the time being, what is important to note is that the model of attitudes as object-evaluation associations of varying strength absorbs the earlier-mentioned "lessons." It does not presuppose attitude-behavior consistency as part of its very definition. Thus, it avoids DeFleur and Westie's "fallacy of expected correspondence." The model also is clear about the status of attitude as evaluative knowledge, akin to any other knowledge in terms of its representation in memory. Yet, the model specifies mechanisms by which this representation can influence behavior and points to the strength of the object-evaluation association in memory as a determinant of the extent to which any such influence is likely. Thus, the model provides a perspective as to why and how attitude-behavior consistency varies.

Although the matter is somewhat of a digression, it will be useful to highlight one additional point before concluding this section. The use of the term object-evaluation *association* is *not* intended to commit or restrict the theoretical perspective to an associative network model of memory. Of course, one can view an attitude as an evaluative tag that is linked to the representation of the attitude object in memory and that has the potential to receive some activation whenever the object itself is activated. And, this certainly was the dominant approach to memory at the time the model was first proposed. However, the theoretical perspective is also completely consistent with the connectionist approaches that have emerged since that time. An attitude can be viewed as a pattern of activation that emerges as a consequence of the connection weights that a system has learned to apply to a set of input units. Indeed, various colleagues and I have considered attitudes and attitude accessibility from the perspective of connectionist models (Smith, Fazio, & Cejka, 1996; Fazio, Eiser, & Shook, 2004). Moreover,

we have proposed what is one of the most fully detailed connectionist models of attitude development available in the literature and have conducted numerous simulations employing the network architecture (Eiser, Fazio, Stafford, & Prescott, 2003; Eiser, Stafford, & Fazio, 2007). Whereas an associative network approach views the attitude as a discrete, symbolic unit that is stored and capable of activation, the connectionist approach views the attitude as a pattern of activation that is generated by the learned connection strengths between units. What is important from the present perspective is that both approaches allow for an evaluation to be evoked when an object is encountered and both maintain that the likelihood of such attitude evocation is a function of past learning. Indeed, it is the prior learning that constrains a connectionist network to reliably produce roughly the same pattern of activation in response to two similar sets of input units.

(Some) attitudes are represented in memory

Whether one pursues an associative network or a connectionist approach, there exists an abundance of empirical evidence that attitudes are represented in memory and capable of automatic activation upon presentation of the attitude object. Briefly reviewing such evidence, before turning to a direct discussion of some issues that have arisen in the last few years regarding the attitude concept, will be valuable. The evidence I wish to highlight has one central, and I believe compelling, feature. During a task or in a situation in which individuals are presented with attitude objects but provided with no reason to actively consider their attitudes toward the objects, their performance is affected by their attitudes. Thus, presentation of those objects must have evoked the attitudes despite their irrelevance to the immediate task demands.

The most well-known of such research paradigms is the evaluative priming procedure in which attitude objects are presented as primes unrelated to the participants' primary task of identifying the connotation of target adjective (Fazio et al., 1986). Responding is facilitated by the presentation of a prime that is evaluatively congruent with the target adjective. This is not the time to review what is now a very extensive literature (see Fazio, 2001; Klauer & Musch, 2003, for reviews). Instead, I will simply highlight a few findings that are especially relevant to the argument that some pre-existing attitudinal representation is evoked automatically and that the likelihood of such activation occurring is a function of the strength of the object-evaluation association.

Early research concerning the attitude-nonattitude continuum was premised on the notion that, when directly queried about their attitudes, individuals with stronger associations would be able to respond more quickly than those with weaker or absent associations. Thus, latency of response to a direct attitudinal query was employed as a dependent measure assessing such attitude accessibility. These experiments succeeded in identifying a number of variables that decreased the time individuals required to report their attitudes, including direct experience with the attitude object (Fazio et al., 1982), inference from free-chosen as opposed to manded behavior (Fazio, Herr, & Olney, 1984), and most relevantly, repeated attitudinal expression. Noting and rehearsing the association in an early phase of the experiment enhanced the speed with which individuals could later indicate their attitudes (Fazio et al., 1982; Powell & Fazio, 1984). Numerous later experiments have employed both this latency measure of associative strength and the attitude rehearsal manipulation within the evaluative priming procedure. Priming effects have been stronger for attitude objects involving stronger object-evaluation associations (see Fazio, 2001, for a review).

Just as object-evaluation associations can be strengthened by rehearsal, they can be weakened by repeated nonuse. In a recent series of experiments (Sanbonmatsu, Posavac, Vanous, Ho, & Fazio, 2007), attitude objects from one of two sets that involved strong evaluative associations (e.g., holiday, disease) were presented repeatedly (e.g., 40 presentations of each object name)

in a task in which the automatically-activated attitude was irrelevant. Participants were required only to recite the object word aloud as it gradually emerged from a rectangular block of masking dots. Presumably, attitudes were automatically activated upon presentation of the objects, at least early in the task, but they were not of any use. When participants later underwent a priming procedure that included these words, as well as the set of control objects, they displayed less evidence of attitude activation in response to the words that had been repeatedly presented in the earlier task. Such consequences of repeated nonutilization of the attitude were observed even when the two phases of the experiment were separated by a 24-hour delay. Thus, experimental manipulations that serve to strengthen, or to weaken, object-evaluation associations have a corresponding effect on the likelihood of automatic attitude activation.

A conceptually parallel finding regarding the moderating role of attitude strength has been observed at the level of an individual difference. Hermans, De Houwer, and Eelen (2001) compared individuals who had scored extremely high or extremely low on Jarvis and Petty's (1996) Need to Evaluate Scale. The scale assesses the extent to which individuals report having attitudes and evaluating objects regularly. NES moderated the magnitude of the typical evaluative priming effect, with the higher NES participants displaying greater evidence of automatic attitude activation.

Importantly, evidence that at least some attitudes are represented in memory and can be evoked automatically extends well beyond the evaluative priming paradigm. A variety of findings point to an influence of attitudes on tasks during which individuals experience no need to consider their evaluations of the presented objects. For example, Roskos-Ewoldsen and Fazio (1992) observed effects of attitude accessibility on a variety of tasks related to visual attention. More attitude-evoking objects – whether measured via latency of response to a query or experimentally manipulated – attracted more attention. When presented as distractors, attitude-evoking objects were more likely to be noticed incidentally in one such experiment, and more likely to interfere with performance of a visual search task in another.

In research concerning the construal of multiply-categorizable objects, Smith et al. (1996) demonstrated that categories made more attitude-evoking by virtue of an attitude rehearsal manipulation were more influential. The more attitude-evoking categories were more likely to govern interpretation of the multiply-categorizable objects. When the objects (e.g., sunbathing, Pete Rose) were presented as cues in a cued-recall test, they more effectively cued the categories for which object-evaluation associations had been strengthened than categories from a control set (e.g., cancer versus beach, and baseball player versus gambler, depending upon which had been assigned to the attitude rehearsal task versus the control task). In one experiment, this effect was observed despite the imposition of a delay of one full week between the attitude rehearsal manipulation and the cued recall test. Recent research by Ferguson, Bargh, and Nayak (2005) has yielded related evidence regarding the impact of automatically-activated attitudes on construals of subsequently presented stimuli. Although associative strength was not included as a factor in the design, positively- and negatively-valued attitude objects were selected as stimuli on the basis of normative data showing them to involve relatively strong object-evaluation associations. Subliminal presentation of these object names influenced how subsequent ambiguous stimuli were categorized. The ambiguous objects were more likely to be classified as relating to a more positive category when they had been preceded subliminally by a positively-valued object.

Conceptually parallel effects have been observed with respect to photos of multiply-categorizable persons. In a two-session experiment, Fazio and Dunton (1997) first assessed participants' automatically-activated racial attitudes via an evaluative priming procedure. As is typical in such research, a full range of attitude estimates were obtained. Some participants showed evidence of automatically-activated negativity in response to Black faces; for others,

positivity was evoked; and for others, neither occurred. In the second session, participants made similarity judgments of pairs of stimulus persons who varied in race, gender, and occupation. Those similarity judgments were more strongly influenced by race for those participants for whom race was more attitude-evoking, in either the positive or the negative direction. In other words, a curvilinear, U-shaped relation was observed between the estimates of automatically-activated racial attitudes and categorization by race. Apparently, individuals for whom race was attitude-evoking were more likely to have their attention drawn to the target's skin color and, as a result, to use race (as opposed to gender or occupation) more heavily as a basis for judging similarity.

As an additional example of research demonstrating that attitudes can be evoked and exert an influence in situations in which they are not relevant to task demands, we can consider experiments that have examined the effects of accessible attitudes on sensitivity to changes in the appearance of the attitude object (Fazio, Ledbetter, & Towles-Schwen, 2000). This work involved "head-and-shoulders" photos of undergraduates. At the beginning of the experiment, participants were exposed to these photos multiple times and each time expressed their attitudes (how attractive they personally found the photo) or engaged in a control task (estimating the individual's height). A subsequent task involved the presentation of these original photos, as well as computer-generated morphs representing varying degrees of change from the original photos. Participants were to note whether any given photo was exactly the same as one presented earlier or different in any way. The two conditions did not differ, in either error rates or latencies, when judging the unchanged stimuli, thus establishing that the original photos were encoded equally well. However, the conditions did differ with respect to the morphs, especially with respect to those that more closely resembled the originals and, hence, were likely to evoke the previously rehearsed attitude. The attitude rehearsal participants had relatively more difficulty detecting the presence of change (either more errors or longer latencies to answer correctly, across the various experiments). They also perceived less change when they did detect it, viewing the morph as relatively more likely to be a different photo of a person seen before than a photo of a novel person never before seen. Thus, stimuli that were changed, but remained sufficiently similar to the original to evoke the attitude associated with the original stimulus, were assimilated in the direction of the attitude. This assimilation apparently contributed to an initial sense of familiarity with the changed stimuli, making the detection of change relatively more difficult.

Finally, research by Chartrand and Bargh (1999, as discussed in Bargh & Chartrand, 1999) has examined the interesting possibility that differential mood states can sometimes arise, for seemingly unknown reasons, simply as a function of repeated exposure to valenced stimuli. In the context of an alleged visual acuity task that involved indicating as quickly as possible the side of the screen on which flashes appeared, participants were subliminally exposed to four nouns numerous times. The between-subjects design involved nouns that normative data had shown to be strongly positive (e.g., music, friends), strongly negative (cancer, cockroach), weakly positive (parade, clown), weakly negative (Monday, worm), or neutral (building, plant). Relative to the neutral condition, subliminal exposure to the objects involving strong evaluative associations, but not those involving weaker associations, influenced subsequent mood reports. Thus, repeated automatic activation of positivity (or negativity), as a result of the exposure to objects involving strongly associated positive (or negative) evaluations, generated a positive (or negative) mood.

I may have belabored the point, but these various findings all illustrate an influence of attitudes, especially ones involving stronger object-evaluation associations, on subsequent task performance or judgments. Importantly, this influence is evident even though the subsequent situation provides no reason for actively considering those attitudes. Thus, visual search performance, cued recall, judgments of similarity, perceptual assimilation, and reports of

current mood were all affected by the presence (sometimes subliminal presence) of attitude-evoking objects. The findings converge on one fundamental conclusion. At least some attitudes are represented in memory, and they are activated automatically when the attitude objects are encountered.

(Some) attitudes are constructed

As emphasized by the notion of object-evaluation associations varying in strength and the attitude-nonattitude continuum, not all objects are represented in memory with strongly associated evaluative linkages. Yet, situational demands sometimes force individuals to make evaluative judgments and decisions regarding novel entities (e.g., the prospect of dining at a newly-opened restaurant or voting for a newly-emerging candidate for public office) or alternatives within what is typically a domain of indifference (e.g., persons who are barely attentive to college basketball during the regular season, but annually enter brackets into an NCAA tournament pool). In such cases, attitudes need to be constructed. The effort and deliberation devoted to such construction will vary as a function of just how motivated the individual might be to reach an accurate assessment (Kruglanski, 1989; Fazio, 1990; Tetlock & Kim, 1987). An evaluation may be constructed relatively effortlessly on the basis of sheer resemblance to entities for which one already has attitudes represented in memory (e.g., Duckworth, Bargh, Garcia, & Chaiken, 2002; Gilovich, 1981; Shook, Fazio, & Eiser, in press). Or, the individual may deliberate extensively about the specific attributes that characterize the entity and their favorability (e.g., Sanbonmatsu & Fazio, 1990).

Because attitudes vary in strength, they are differentially sensitive to information that is salient in the immediate situation. Such a theoretical assertion is by no means novel. Indeed, Bem's (1972) impactful proposal that "individuals come to 'know' their own attitudes, emotions, and other internal states partially be inferring them from observation of their own behavior and/or the circumstances in which this behavior occurs" was accompanied by the proviso that such attitudinal inference processes occur "to the extent that internal cues are weak, ambiguous, or uninterpretable" (p. 2). A considerable amount of evidence has accumulated in support of this reasoning. The relevant research typically assumes a particular form. A manipulation makes salient either positive or negative information related to the attitude object, and the impact of that manipulation on subsequent attitudinal reports is compared in two samples that differ with respect to some index of the strength of the participants' initial attitudes. For example, Chaiken and Baldwin (1981) employed a linguistic manipulation, based on earlier research by Salancik and Conway (1975) to make salient to individuals either pro-ecology or anti-ecology behaviors that they had performed. Due to how the adverbs "occasionally" and "frequently" were paired with various ecology-related behaviors, participants found themselves indicating that they had performed many pro and few anti behaviors, or vice versa. This linguistically-biased review of their past behaviors led participants to differential views of their favorability toward environmentalism. However, it did so more for some individuals than for others. In an earlier session, participants had been administered various measures of attitude – a cognitively-oriented measure based on an expectancy-value framework and a more affectively-oriented measure. Only individuals who had responded relatively inconsistently across the two different measures, suggesting that they lacked clear knowledge of their attitudes, were influenced by the linguistically-biased recall of past behaviors.

Similarly, Hodges and Wilson (1993) examined the moderating role of attitude accessibility on the effect that analyzing reasons has on attitudinal reports. Immediately after explaining the reasons for their attitudes, people report attitudes that are very consistent with the valence implied by the now very salient reasons that they articulated. However, these easily verbalizable reasons may be an unrepresentative subset and, hence, may be somewhat discrepant with the individuals' attitudes and with their prior reports (see Wilson, Hodges, &

LaFleur, 1995). Hodges and Wilson (1993) found that the extent to which such consequences occurred varied as a function of the accessibility of individuals' attitudes. Those who responded relatively quickly to an attitudinal query at Time 1 were less influenced by verbalizing reasons for their attitudes at a later time than were those who responded slowly. The latter group – those with lower attitude accessibility – displayed a stronger correlation between their Time 2 attitude reports and the favorability implied by the reasons they articulated and, as a result, a weaker correlation between their Time 1 and Time 2 reports. These differences occurred despite the fact that the two groups displayed equivalence in terms of the extent to which the reasons they generated were consistent with the Time 1 ratings. Thus, even though they verbalized equivalently unrepresentative reasons, individuals with greater attitude accessibility were less influenced by these momentarily salient reasons when they were reporting their attitudes at Time 2, and their reports displayed greater stability across time.

Relatedly, many additional studies have found the stability of attitude reports over time and/or resistance to counterattitudinal information to vary as a function of attitude accessibility (e.g., Bassili, 1996; Bassili & Fletcher, 1991; Fazio & Williams, 1986; Zanna, Fazio, & Ross, 1994). Moderating effects of the sort reviewed in this section clearly illustrate the importance of attitude strength. Contextually salient information can influence the attitudinal ratings that people provide at any given moment in time, but the extent of that influence varies. When objects are represented in memory with relatively weakly associated (or absent) evaluative linkages, individuals faced with a query, or some other situational demand for evaluative judgment, will need to construct an attitude. Hence, they will be more affected by the currently salient information than individuals for whom a preexisting evaluative association is activated automatically.

Problems with a Strong "Attitudes as Constructions" Perspective

In light of the research summarized in the preceding section, I have to admit to some bewilderment regarding recent discussions portraying all (or virtually all) attitudes as momentary constructions (e.g., Schwarz & Bohner, 2001; Zaller & Feldman, 1992). These formulations view attitudes as evaluative judgments that are always computed from scratch on the basis of information that is accessible at that moment, as opposed to their being represented and activated from memory. Thus, the strong version of this perspective is contrary to the evidence that at least some attitudes are represented in memory and exert influence even when individuals have no reason to reflect upon, retrieve, or construct evaluations.

The argument typically advanced as the reason to entertain a constructionist perspective centers on the contextual dependence of evaluative judgments. There can be no question that attitudinal reports can be influenced by salient contextual information. We already have highlighted some such evidence, for example, Chaiken and Baldwin (1981) and Hodges and Wilson (1993). Yet, this research demonstrates not only contextual dependence, but also moderation of this dependence as a function of attitude strength.

In my view, then, the strong constructionist perspective is contradicted by existing data. In addition, as Fazio and Olson (2003a) argued, it is characterized by an inherent logical problem. To illustrate, let's consider an individual confronted with a situational demand that requires evaluation of attitude object X. Because X is novel (e.g., a new automobile model or a political candidate) and bears no obvious resemblance to objects toward which attitudinal associations exist in memory, the individual will need to engage in some degree of construction. Presumably, the individual will compute an attitude toward X on the basis of whatever attributes of X happen to be salient at the time. Let's presume that these include attributes A, B, and C. If these attributes are themselves positively-valued, then the individual will conclude that X is positive. If they are predominately negative, then X will be viewed as negative. But,

from where did the evaluations of attributes A, B, and C emerge? Why did they take on positive or negative value? The strict constructionist perspective has to maintain that the attitudes toward the attributes were themselves the outcome of a construction process. So, salient attributes of A – a_1, a_2, and a_3 – inform the computation of the attitude toward A. But, why does a_1 assume a particular valence? Obviously, the construction process is open to a problem of infinite regress. Ultimately, some relevant evaluation – some outcome of prior learning history – needs to be represented in memory.

The importance of prior learning points to yet another conceptual problem with the strong constructionist perspective. What kind of information processing system would not learn from its computational efforts? Would any functional system, once it has concluded that X is good, lose all trace of its construction output? Would it have to "re-compute" entirely from scratch with no savings in relearning the next time a situation demands an evaluative judgment of X? If attitudes are viewed as evaluative knowledge, akin to any other knowledge, why wouldn't such knowledge accumulate over time? Associative models of learning and memory call for the development or strengthening of an excitatory evaluative link. Connectionist models call for a change in the learned weights connecting the input units and the output.

Social and cognitive scientists do not question that people readily acquire semantic knowledge. They learn that the probability of encountering "butter" is greater given the presence of "bread" than its absence, that the probability of seeing a "nurse" increases in the presence of a "doctor." Research on impression formation and social perception has repeatedly demonstrated that perceivers develop associations between a target person and traits that characterize the person's behavior (e.g., Carlston & Skowronski, 2005; Uleman, 1999). Why should evaluative knowledge be any different?

Even concepts that at first require extensive computational resources clearly benefit from learning. A young child performs basic arithmetic operations such as addition by counting. So, to determine the total number of balls that are available, the child points and counts the two red balls, "one, two," and continues counting "three, four" while pointing to the two blue balls. Knowing the principle of cardinality, the child proudly announces that there are "four" balls. The answer is obviously the outcome of a construction process. But over time and with practice, the child needs to rely less and less on pointing, counting, and procedural rules. The child comes to know that 2 + 2 = 4. The output of the previous constructions is learned and represented in memory as the sum of two sets of two. Again, why should evaluative knowledge be any different?

On the basis of these arguments alone, I fail to see how a strong "attitudes as constructions" perspective can be granted credence. Nor am I alone. Even proponents seem to agree: "…we adopt a strong version of a construal model, assuming, for the sake of the argument, that respondents *always* [emphasis in original] need to compute a judgment from scratch and can't recall their previous evaluations. As anyone who remembers that a movie was "boring" – but can't recall any relevant details – realizes, this extreme assumption is unrealistic" (Schwarz & Bohner, 2001, p. 444). It is this very unrealism that necessitates consideration of an *attitude-nonattitude continuum*. Attitudes are evaluative knowledge represented in memory, but there will be objects for which individuals lack such knowledge and, hence, need to construct overall evaluations on the basis of whatever relevant information they can access from memory or the situation makes available.

However, the outcomes of any such construction efforts are not forgotten. To the contrary, they facilitate progression from the status of a nonattitude toward the status of a strongly associated evaluation that can be activated automatically when the object is encountered. This attitude development has functional value for the individual, enabling efficient and relatively effortless

appraisal of the object. For example, attitude rehearsal has been shown to ease later decision-making, decreasing the resources that are needed to reach evaluative decisions. When asked to make rapid decisions about their preferences between pairs of objects, individuals who had earlier been induced to develop and rehearse attitudes toward novel objects displayed less cardiovascular reactivity than individuals who had been exposed to the objects equally often but in the context of a task that was unrelated to attitudes (Blascovich, et al., 1993; Fazio, Blascovich, & Driscoll, 1992). Likewise, less reactivity was exhibited when individuals made decisions involving objects for which they had earlier rehearsed attitudes than objects for which they had not (Blascovich, et al., 1993). Activation of the rehearsed attitudes during the pairwise preference task apparently made the decision task less demanding. Fewer resources were required to cope with the demands of the task (see Fazio, 2000, for an extensive discussion of research related to the object appraisal function of attitudes). Thus, position along the attitude-nonattitude continuum influences the very nature of the evaluative judgment process.

The Model's Perspective Regarding a Few Contemporary Issues

Responses to explicit measures of attitude are verbal behaviors

I have tried to be careful and precise in this paper about my use of the term attitude versus individuals' reports of their attitudes. Responses to a verbal query should not be viewed as attitudes per se. They are verbal expressions of the attitude, subject to many forces beyond the influence of any evaluative associations that might exist in memory. The distinction is one that attitude theorists have long recognized. In his classic article "Attitudes can be measured," Thurstone (1928) used the term "opinion" to refer to "a verbal expression of attitude" and argued that an opinion merely "symbolizes an attitude" (p. 531). Although Thurstone advocated using "opinions as the means for measuring attitudes," he noted that "neither [a person's] opinions nor his overt acts constitute in any sense an infallible guide to the subjective inclinations and preferences that constitute his attitude. Therefore, we must remain content to use opinions, or other forms of action, merely as indices of attitude" (p.532).

Thurstone was concerned that a person's verbal behavior may be influenced by factors in addition to the attitude. He specifically made reference to possibilities ranging from intentional misrepresentation in the interest of avoiding outright dispute with others to less than frank expressions stemming from courtesy. But, there are many factors in addition to self-presentation concerns that can make responses to verbal queries poor indices of the attitude of interest. As many survey researchers have noted (e.g., Schuman & Presser, 1981; Schwarz & Bohner, 2001; Tourangeau & Rasinksi, 1988), even well-intentioned responding to a verbal measure of attitude is a complex process involving issues of question comprehension, scale interpretation, and the use of appropriate standards of comparison. In this sense, verbal *reports* of attitudes are always constructions, and those responses can be influenced by construals and standards that bear little relation to what the survey researcher intended.

Correspondence between implicit and explicit measures—Even if the "communication" between the questioner and the respondent has proceeded as intended, the verbal report may be influenced by many factors in addition to the evaluative association represented in memory. By not requiring direct reports, implicit measures of attitude aim to provide a more proximal estimate of these attitudinal representations than is possible with explicit measures. This difference alone may result in an observed discordance between the attitude estimates provided by implicit and explicit measures. My colleagues and I have approached this issue of correspondence between implicit and explicit measures from the perspective of the MODE model of attitude-behavior processes (Fazio, 1990; Fazio & Towles-Schwen, 1999; Fazio & Olson, 2003b). Because responding to an explicit measure is itself a verbal behavior, it can be viewed as related to the multiple attitude-behavior processes that form the focus of the MODE model. The verbal expression may stem from a relatively

spontaneous process involving the direct influence of the automatically-activated attitude on the immediate perceptions of the attitude object noted in the question. Alternatively, it may stem from a very deliberative analysis of the attributes that characterize the attitude object, provided the individual has both the motivation and the opportunity to engage in such effortful reflection. Finally, the verbal response may stem from what the MODE model refers to as mixed processes, involving both automatic and controlled components. In such cases, the verbal expression may be influenced not only by the relevant automatically-activated attitude, but also by the downstream consequences of any number of motivational factors. Individuals may be so concerned about accuracy that they check the implications of the activated attitude against their knowledge regarding the attributes of the object, instead of just "going with their guts." Or, sufficiently motivated individuals may compare the activated attitude for its consistency with various internal standards and cherished values. And, of course, as Thurstone noted, they may "spin" their attitudes so as to make their verbal responses more palatable to the apparent audience.

According to the MODE model, then, any attitudes activated automatically by the mention of the attitude object in an explicit measure are merely a "starting point" for the verbal behavior. Whether the verbal response reflects the attitude will depend on the motivation and opportunity factors. Verbally expressed judgments occur farther "downstream" than the automatic activation of any relevant attitude and may be influenced by motivational factors that can override the effect of the automatically-activated attitude. For these reasons, greater correspondence between implicit and explicit measures is expected when the motivation to deliberate and/or the opportunity to do so are low. When both are high, any additional information that the motivated individual considers has the potential to produce discordance between the verbal responses and the attitude estimated by the implicit measure. This reasoning has received empirical support in a number of investigations concerned with race-related judgments (see Fazio & Olson, 2003b, for a review).

Neither unconscious nor dual attitudes—Importantly, from the perspective of the MODE model, any observed discordance between implicit and explicit measure does not necessitate reference to the unconscious (e.g., Banaji, 2001) or to dual representations of attitude (e.g., Wilson, Lindsey, & Schooler, 2000). As Fazio and Olson (2003b) emphasized, no implicit measure speaks to the question of whether individuals are or are not aware of their attitudes. Additional evidence, well beyond the administration of an implicit measure, is required to make any such inference (see Olson & Fazio, 2002, for an example). In addition, no explicit measure guarantees the existence of a distinct representation in memory, independent of the attitude that is automatically activated when the object is encountered. Responding to the explicit measure is an exercise in verbal behavior and, hence, does not speak to the presence or absence of pre-existing representations.

On the presumed malleability of automatically-activated attitudes

Research involving implicit measures of attitude focused at first on their correspondence with explicit measures and their utility as predictors of judgments and behavior. However, interest in the possibility of change in automatically-activated attitudes, especially prejudice, spawned research in which implicit measures served as dependent measures. The findings (see Blair, 2002, for a detailed review) revealed what have been viewed by some as a surprising sensitivity of implicit measures to immediate contextual factors and, hence, have raised questions about the relative stability versus malleability of automatically-activated attitudes. In my view, the case for malleability has been overstated. What has been observed is malleability of the scores on the implicit measure. I do not regard such findings as necessarily indicative of malleability of the attitude. Instead, the malleability may be more precisely pinpointed as stemming from

either the malleability of the measure or the malleability of the attitude object. It is important that the distinctions not be blurred.

Malleability of the measure: The IAT's sensitivity to extrapersonal associations—Considerable recent research has demonstrated that pre-exposure to different kinds of information influences IAT scores (see Blair, 2002, for a review). Essentially, salient positive information about a given category produces IAT scores indicative of more favorable evaluations. Such manipulations have involved exposure to movie clips that depicted Blacks at a harmonious family event versus an argumentative, gang-related scene (Wittenbrink, Judd, & Park, 2001), presentation of a series of either admired Black individuals and disliked Whites or disliked Blacks and admired Whites (Dasgupta & Greenwald, 2001), assigning participants to interact with a Black partner who occupied a superior or subordinate task role (Richeson & Ambady, 2003), exposure to violent and misogynous rap music (Rudman & Lee, 2002), and introduction to the IAT (specifically as a procedure for assessing prejudice) by a Black versus a White experimenter (Lowery, Hardin, & Sinclair, 2001). Yet, such findings do not constitute unambiguous evidence for the inference that automatically-activated attitudes have been modified. The critical issue in need of resolution is whether such contextual manipulations (a) produce an actual change in the relevant mental representation, to which the IAT is sensitive, or (b) simply render momentarily salient information that assists the participant with the response mapping problem posed by the IAT.

The findings that my colleagues and I have recently accumulated regarding the sensitivity of the IAT to extrapersonal associations lead me to favor the second alternative, or at least to the belief that considerable caution is required in drawing any inferences about attitude change on the basis of such IAT findings. Extrapersonal associations – ones that neither form the basis of the individual's attitude nor become activated automatically in response to the object – can influence the IAT as it is traditionally implemented. Especially relevant is research in which both attitudes and extrapersonal associations were experimentally created in the lab (Han, Olson, & Fazio, 2006). Participants formed a preference for one novel stimulus over another, with the former being described in ways that made it objectively far superior. They then were exposed to videotaped comments of two young boys who expressed opinions that were either consistent with that preference or incredulously inconsistent. The results revealed that those given attitudinally-inconsistent extrapersonal associations were able to use that information to facilitate coping with the incompatible mapping problem posed by the traditional IAT. They exhibited a significant reduction in their presumed preference for the objectively superior of the two stimuli, compared to those who had received the consistent extrapersonal association. This attenuation of the IAT effect occurred despite the fact that the participants rated the sources of the attitude-inconsistent extrapersonal information as irrational and foolish. On the other hand, the extrapersonal associations did not influence either a subliminal priming measure, or a version of the IAT that was "personalized" and, hence, less susceptible to the influence of extrapersonal associations (Olson & Fazio, 2004). These measures proved sensitive to the attitude, regardless of the congruency of the extrapersonal information. Thus, the various findings noted earlier regarding contextual sensitivity may reflect the influence of momentarily salient extrapersonal associations. As such, they may point more to the malleability of the IAT than to the malleability of attitudes.

Malleability of the attitude object: Let's not forget Asch—Some of the research illustrating contextual sensitivity has employed, not the IAT, but an evaluative priming procedure. When Black faces are presented within the background context of a church interior, they prime less negativity than when the context is an urban street corner or a jail (Wittenbrink et al., 2001; Barden, Maddux, Petty, & Brewer, 2004). Similarly, current motivational states have been shown to influence the evaluation that is primed by the presentation of a related object. Water primes more positivity when individuals are thirsty (Ferguson & Bargh, 2004),

food stimuli activate more positivity when individuals are hungry (Seibt, Hafner, & Deutsch, 2007), and cigarettes prime more positivity when heavy smokers are in a state of deprivation as opposed to having just had a cigarette (Sherman, Rose, Koch, Presson, & Chassin, 2003). Whether the context is internal or external to the individual, findings of this sort need not be interpreted as evidence of the malleability of automatically-activated attitudes or of objects being associated with multiple attitudes. Instead, they may represent varying construals of the attitude object.

This is the very issue that Asch (1948) addressed long ago when he argued against the "doctrine of suggestion, prestige, and imitation" as the explanation for the often-observed finding that a statement is evaluated more positively when it is attributed to a more prestigious source. He noted that the finding need not presuppose that "an unchanged object of judgment undergoes change of evaluation" (p. 255) as a function of the prestige of the individual from whom it emanates. The apparent changes in evaluation may be due to "a change in the object of judgment rather than in the judgment of the object" (p. 256). To use the classic example, "rebellion" does not refer to the same entity when attributed to Jefferson versus Lenin. Thus, the statement that "a little rebellion, now and then, is a good thing and necessary in the political world as storms are in the physical" is evaluated differently due to a change of meaning. The importance of such change of meaning has been documented repeatedly in research concerning impression formation (e.g., Hamilton & Zanna, 1974; Higgins & Rholes, 1976).

In parallel fashion, I would argue that water is not the same attitude object when one is thirsty as when not, a given food is not the same object when one is hungry versus sated, and cigarettes are not the same attitude object when a smoker is in a state of craving versus not. Likewise, situational cues can lead to varying construals of any given object. The very same African-American male will conjure very different interpretations depending on whether he is dressed in a three-piece suit versus gang attire. Virtually all persons, objects, and issues are multiply categorizable. It is how an object is construed at any moment in time that determines attitude activation (see Barden et al., 2004; Fazio, 1986; Fazio & Dunton, 1997; Lord & Lepper, 1999; Sinclair & Kunda, 1999; Smith et al., 1996, for discussion of such categorization processes). Thus, it is not surprising that the attitudes activated in a priming experiment may vary as a function of the effect of current concerns and context on construal of the objects.

I do not mean to imply that automatically-activated attitudes are incapable of change. To the contrary, I believe that some research findings provide evidence strongly suggesting that appropriately-designed social influence techniques can produce an actual change in the representation of the attitude (e.g., Kawakami, Dovidio, Moll, Hermsen, & Russin, 2000; Olson & Fazio, 2006; Rudman, Ashmore, & Gary, 2001). However, given the abundance of evidence regarding the relative stability of accessible attitudes over time and their resistance to counterinformation (see Fazio, 1995, for a review), I am skeptical that the mere contextual salience of counterattitudinal information has any impact on the attitude representation. To me, it seems far more likely that any such empirical demonstration reflects the malleability of the implicit measure or the malleability of the object than the malleability of the attitude itself. In any case, I would urge that researchers recognize these multiple possibilities and exercise caution in the inferences that are drawn from the observation of change in implicit measure scores.

Attitude Ambivalence

It is rare for any given object not to be characterized both by some attributes that might be considered positive and by some that might be considered negative. This creates the potential for the experience of ambivalence when judging the object or making an approach/avoidance decision. To the extent that the individual lacks an accessible attitude and is forced to construct an evaluation on the spot in response to a verbal query or a need for action, the individual may

experience ambivalence if both the positive and the negative features of the object are salient (Newby-Clark, McGregor, & Zanna, 2002). However, the essence of more reasoned forms of attitude development is an integration of any such conflicting information into a summary evaluation. Once formed, and associated with the attitude object, this summary evaluation effectively resolves the ambivalence. In future similar situations, the object will evoke the summary evaluation, instead of requiring continued consideration of the attribute information. Thus, ambivalence amounts to pre-decisional conflict and can be obviated by the activation of a relevant summary evaluation developed at an earlier point in time. According to the MODE model, individuals may re-visit the attribute level, if they are sufficiently motivated and have the opportunity to do so (Fabrigar, Petty, Smith, & Crites, 2006; Sanbonmatsu & Fazio, 1990). However, when either motivation or opportunity is lacking, the earlier formed summary evaluation will predominate.

It is important to note that the potential for ambivalent reactions to an attitude object relates very closely to the fact that most objects can be construed in multiple ways. Ambivalence is often portrayed, as I have framed it above, as stemming from the subordinate attribute level (e.g, Kaplan, 1972). One is torn by the object's obvious possession of both positive and negative features (e.g., chocolate cake is both delicious and fattening). However, ambivalence also can be viewed as an issue of categorization (see Fazio, 1994, for an earlier discussion of this perspective). To which category does the object belong? This more superordinate emphasis highlights that the object can be construed in multiple ways. In many situations, contextual cues and current motivations may foster a particular categorization. In the context of a self-indulgent dinner celebrating a significant life event, the chocolate cake is likely to be categorized as a wonderfully tasty dessert and evoke nothing but positivity. Shortly before or after a visit to the gym, the same cake may be quickly relegated to the fattening category and evoke negativity. When contextual cues are less dictating of the categorization outcome, one construal may still be more likely to predominate over the other as a function of the individual's own proclivities. As noted earlier while discussing the influence of attitude accessibility on categorization, the more attitude-evoking possibility is at an advantage (Fazio & Dunton, 1997; Smith et al., 1996.) Thus, committed and successful dieters are likely to have developed such strongly associated negativity to fattening foods that they are more apt to "assign" the chocolate cake to the category of "fattening" than are individuals who have the advantage of a high metabolism (see Ferguson, 2007; Fishbach, Friedman, & Kruglanski, 2003, for related research findings). Neither of these two types of individuals is likely to experience much ambivalence. Ambivalence is more likely to arise in situations in which, and among people for whom, potential categorizations essentially compete for attention. Viewed from this perspective of superordinate categorization, the ambivalence stems, not from the simultaneous activation of positive and negative sentiments, but from shifting (possibly rapidly shifting) construals of the attitude object. When categorized in one way, the chocolate cake evokes positivity. When categorized in another way, it activates a negative attitude. Thus, it is the malleability of the attitude object, not the malleability of the attitude or the possession of multiple attitudes toward the object, that dictates the experience of ambivalence.

Such a view has important implications for the measurement of attitudes and attitude ambivalence. When we as researchers ask individuals for their evaluations of an object, we often do so in a rather abstract manner that lacks specific contextual cues. As a result, we may be requesting evaluation at a more superordinate level of categorization than respondents typically employ for that object or set of objects. Individuals' awareness of the potential for multiple categorizations and their consideration of varying evaluations may prompt them to express ratings indicative of ambivalence, effectively noting that their evaluations "depend" on the details. Yet, in a specific situation, any such ambivalence may not be experienced because sufficient cues exist to disambiguate the categorization and, hence, prompt activation of the attitude associated with a single category.

Summary and Conclusions

The major message that I hope this article has succeeded in communicating can be summarized in two words. Attitudes exist! They are not a scientific reification. They are not hypothetical constructs. Attitudes are evaluative knowledge, represented in memory in the same way as any other form of knowledge. They summarize our prior learning with respect to the valence of the outcomes produced by a given object. As summary evaluations associated with the representation of an attitude object, these attitudes can be activated from memory automatically when the object (or a sufficiently related object) is encountered. According to the MODE model, such automatically-activated attitudes serve as the "starting point" for our appraisals of the object in the immediate situation. When either the motivation to deliberate further, or the opportunity to do so, is lacking, these immediate perceptions will influence evaluative judgments, verbal expressions, decisions, and overt behavior in a relatively direct manner. In this way, the evaluative knowledge that we have acquired as a result of our prior learning history – our collection of attitudes – proves to be very functional. Attitudes simplify our day-to-day existence, enabling efficient appraisal of the objects that we encounter. Although relatively thoughtless, these appraisals promote approach behavior toward objects from which we probabilistically will experience pleasure and avoidance of objects likely to produce pain. Thus, attitudes form the cornerstone of a truly functional system by which learning and memory guide behavior in a fruitful direction.

It is important to recognize, however, that the direct "downstream" consequences of automatically-activated attitudes are not inevitable. The automatically-activated evaluation is a starting point or default value, but its influence can be overridden. Under some circumstances (the presence of both motivation and opportunity), individuals may check the validity of their immediate appraisals against other knowledge they have regarding the attitude object or examine its consistency with motivational goals to which they aspire. Such knowledge and motivational forces can attenuate the impact of the automatically-activated attitude or even produce an effort to correct for its influence.

Responses to explicit measures of attitudes are themselves verbal behaviors. They are expressions of the attitude, not the attitude per se. As with any behavior, these verbal expressions can be influenced by not only any attitude that is automatically-activated but also downstream motivational forces. Hence, discordance between implicit measures assessing the automatically-activated attitude and explicit measures can occur, especially under conditions characterized by a motivation to deliberate and opportunity to do so. The discordance is readily understandable in terms of motivated correction for the influence of the automatically-activated attitude. In and of itself, then, any such discordance does not require reference to the unconscious or to dual representations of attitudes in memory.

Acknowledgements

Preparation of this article was supported by Grant MH38832 from the National Institute of Mental Health. The author thanks William Cunningham, Anna Han, Valerie Jefferis, Christopher Jones, and Natalie Shook for their helpful comments on an earlier draft.

References

Allport, GW. Attitudes. In: Murchison, C., editor. Handbook of social psychology. Worcester, Mass.: Clark University Press; 1935. p. 798-844.

Asch SE. The doctrine of suggestion, prestige, and imitation in social psychology. Psychological Review 1948;55:250–276. [PubMed: 18887584]

Banaji, MR. Implicit attitudes can be measured. In: Roediger, HL., III; Nairne, JS.; Neath, I.; Surprenant, A., editors. The nature of remembering: Essays in honor of Robert G. Crowder. Washington, DC: American Psychological Association; 2001. p. 117-150.

Bargh, JA. The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In: Wyer, RS.; Srull, TK., editors. Handbook of social cognition. Vol. 1. Hillsdale, NJ: Erlbaum; 1994. p. 1-40.

Bargh JA, Chartrand TL. The unbearable automaticity of being. American Psychologist 1999;54:462–479.

Barden J, Maddux WW, Petty RE, Brewer MB. Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. Journal of Personality and Social Psychology 2004;87:5–22. [PubMed: 15250789]

Bassili JN. Meta-judgmental versus operative indexes of psychological attributes : The case of measure of attitude strength. Journal of Personality and Social Psychology 1996;71:637–653.

Bassili JN, Fletcher JF. Response-time measurement in survey research: A method for CATI and a new look at nonattitudes. Public Opinion Quarterly 1991;55:331–346.

Bem, DJ. Self-perception theory. In: Berkowitz, L., editor. Advances in experimental social psychology. Vol. 6. New York: Academic Press; 1972. p. 1-62.

Blair I. The malleability of automatic stereotypes and prejudice. Personality and Social Psychology Review 2002;6:242–260.

Blascovich J, Ernst JM, Tomaka J, Kelsey RM, Salomon KL, Fazio RH. Attitude accessibility as a moderator of autonomic reactivity during decision making. Journal of Personality and Social Psychology 1993;64:165–176. [PubMed: 8433270]

Briñol P, Petty RE, Tormala ZL. The malleable meaning of subjective ease. Psychological Science 2006;17:200–206. [PubMed: 16507059]

Bruner JS. On perceptual readiness. Psychological Review 1957;64:123–152. [PubMed: 13420288]

Carlston DE, Skowronski J. Linking versus thinking: Evidence for the different associative and attributional bases of spontaneous trait transference and spontaneous trait inference. Journal of Personality and Social Psychology 2005;89:884–898. [PubMed: 16393022]

Chaiken S, Baldwin MW. Affective-cognitive consistency and the effect of salient behavioral information on the self-perception of attitudes. Journal of Personality and Social Psychology 1981;41:1–12.

Converse, PE. Attitudes and non-attitudes: Continuation of a dialogue. In: Tufte, ER., editor. The quantitative analysis of social problems. Reading, MA: Addison-Wesley; 1970. p. 168-189.

Dasgupta N, Greenwald AG. On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. Journal of Personality and Social Psychology 2001;81:800–814. [PubMed: 11708558]

DeFleur ML, Westie FR. Attitude as a scientific concept. Social Forces 1963;42:17–31.

De Houwer J, Thomas S, Baeyens F. Associative learning of likes and dislikes: A review of 25 years of research on human evaluative conditioning. Psychological Bulletin 2001;127:853–869. [PubMed: 11726074]

Deutsch R, Strack F. Duality models in social psychology: From dual processes to interacting systems. Psychological Inquiry 2006;17:166–172.

Duckworth KL, Bargh JA, Garcia M, Chaiken S. The automatic evaluation of novel stimuli. Psychological Science 2002;13:513–519. [PubMed: 12430834]

Eiser JR, Fazio RH, Stafford T, Prescott TJ. Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations. Personality and Social Psychology Bulletin 2003;29:1221–1235. [PubMed: 15189584]

Eiser, JR.; Stafford, T.; Fazio, RH. How negative attitudes resist change: A connectionist account of the consequences of global and targeted expectancies. University of Sheffield; 2007. Unpublished manuscript

Fabrigar L, Petty RE, Smith SM, Crites SL. Understanding knowledge effects on attitude-behavior consistency: The role of relevance, complexity, and amount of knowledge. Journal of Personality and Social Psychology 2006;90:556–577. [PubMed: 16649855]

Fazio, RH. How do attitudes guide behavior?. In: Sorrentino, RM.; Higgins, ET., editors. The Handbook of motivation and cognition: Foundations of social behavior. New York: Guilford Press; 1986. p. 204-243.

Fazio, RH. Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In: Zanna, MP., editor. Advances in experimental social psychology. Vol. 23. New York: Academic Press; 1990. p. 75-109.

Fazio, RH. Attitudes in associated systems theory. In: Wyer, RS., editor. Advances in Social Cognition. Vol. VII. Hillsdale, NJ: Erlbaum; 1994. p. 157-167.

Fazio, RH. Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In: Petty, RE.; Krosnick, JA., editors. Attitude strength: Antecedents and consequences. Hillsdale, NJ: Erlbaum; 1995. p. 247-282.

Fazio, RH. Accessible attitudes as tools for object appraisal: Their costs and benefits. In: Maio, G.; Olson, J., editors. Why we evaluate: Functions of attitudes. Mahwah, NJ: Erlbaum; 2000. p. 1-36.

Fazio RH. On the automatic activation of associated evaluations: An overview. Cognition and Emotion 2001;15:115–141.

Fazio RH, Blascovich J, Driscoll DM. On the functional value of attitudes: The influence of accessible attitudes upon the ease and quality of decision making. Personality and Social Psychology Bulletin 1992;18:388–401.

Fazio RH, Chen J, McDonel EC, Sherman SJ. Attitude accessibility, attitude-behavior consistency and the strength of the object-evaluation association. Journal of Experimental Social Psychology 1982;18:339–357.

Fazio RH, Dunton BC. Categorization by race: The impact of automatic and controlled components of racial prejudice. Journal of Experimental Social Psychology 1997;33:451–470.

Fazio RH, Eiser JR, Shook NJ. Attitude formation through exploration: Valence asymmetries. Journal of Personality and Social Psychology 2004;87:293–311. [PubMed: 15382981]

Fazio RH, Herr PM, Olney T. Attitude accessibility following a self-perception process. Journal of Personality and Social Psychology 1984;47:277–286. [PubMed: 6481616]

Fazio RH, Ledbetter JE, Towles-Schwen T. On the costs of accessible attitudes: Detecting that the attitude object has changed. Journal of Personality and Social Psychology 2000;78:197–210. [PubMed: 10707329]

Fazio, RH.; Olson, MA. Attitudes: Foundations, functions, and consequences. In: Hogg, MA.; Cooper, J., editors. The Sage Handbook of Social Psychology. London: Sage; 2003a. p. 139-160.

Fazio RH, Olson MA. Implicit measures in social cognition research: Their meaning and use. Annual Review of Psychology 2003b;54:297–327.

Fazio RH, Powell MC, Herr PM. Toward a process model of the attitude-behavior relation: Accessing one's attitude upon mere observation of the attitude object. Journal of Personality and Social Psychology 1983;44:723–735.

Fazio, RH.; Roskos-Ewoldsen, DR. Acting as we feel: When and how attitudes guide behavior. In: Brock, TC.; Green, MC., editors. Persuasion: Psychological insights and perspectives. Thousand Oaks, CA: Sage; 2005. p. 41-62.

Fazio RH, Sanbonmatsu DM, Powell MC, Kardes FR. On the automatic activation of attitudes. Journal of Personality and Social Psychology 1986;50:229–238. [PubMed: 3701576]

Fazio, RH.; Towles-Schwen, T. The MODE model of attitude-behavior processes. In: Chaiken, S.; Trope, Y., editors. Dual process theories in social psychology. New York: Guilford; 1999. p. 97-116.

Fazio RH, Williams CJ. 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 1986;51:505–514. [PubMed: 3761146]

Ferguson MJ. On the automatic evaluation of end-states. Journal of Personality and Social Psychology 2007;92:596–611. [PubMed: 17469947]

Ferguson MJ, Bargh JA. Liking is for doing: The effects of goal pursuit on automatic evaluation. Journal of Personality and Social Psychology 2004;87:557–572. [PubMed: 15535771]

Ferguson MJ, Bargh JA, Nayak DA. After-affects: How automatic evaluations influence the interpretation of subsequent, unrelated stimuli. Journal of Experimental Social Psychology 2005;41:182–191.

Fishbach A, Friedman RS, Kruglanski AW. Leading us not unto temptation: Momentary allurements elicit overriding goal activation. Journal of Personality and Social Psychology 2003;84:296–309. [PubMed: 12585805]

Gawronski B, Bodenhausen GV. Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. Psychological Bulletin 2006;132:692–731. [PubMed: 16910748]

Gilovich T. Seeing the past in the present: The effect of associations to familiar events on judgments and decisions. Journal of Personality and Social Psychology 1981;40:797–808.

Green, BF. Attitude measurerment. In: Lindzey, G., editor. Handbook of Social Psychology. Cambridge, MA: Addison-Wesley; 1954. p. 335-369.

Greenwald, AG.; Nosek, BA. Attitudinal dissociation: What does it mean?. In: Petty, RE.; Fazio, RH.; Briñol, P., editors. Attitudes: Insights from the new implicit measures. Mahwah, NJ: Erlbaum; in press

Hamilton DL, Zanna MP. Context effects in impression formation: Changes in connotative meaning. Journal of Personality and Social Psychology 1974;29:649–654.

Han HA, Olson MA, Fazio RH. The influence of experimentally-created extrapersonal associations on the Implicit Association Test. Journal of Experimental Social Psychology 2006;42:259–272.

Hermans D, De Houwer J, Eelen P. A time course analysis of the affective priming effect. Cognition & Emotion 2001;15:143–165.

Higgins, ET. Knowledge activation: Accessibility, applicability, and salience. In: Higgins, ET.; Kruglanski, AW., editors. Social psychology: Handbook of basic principles. New York: Guilford Press; 1996. p. 133-168.

Higgins ET, Rholes WS. Impression formation and role fulfillment: A “Holistic Reference” approach. Journal of Experimental Social Psychology 1976;12:422–435.

Hilgard ER. Intervening variables, hypothetical constructs, parameters, and constants. The American Journal of Psychology 1958;71:238–246. [PubMed: 13521034]

Hodges SD, Wilson TD. Effects of analyzing reasons on attitude change: The moderating role of attitude accessibility. Social Cognition 1994;11:353–366.

Jarvis WBG, Petty RE. The need to evaluate. Journal of Personality and Social Psychology 1996;70:172–194.

Kaplan KJ. On the ambivalence-indifference problem in attitude theory and measurement: A suggested modification of the semantic differential technique. Psychological Bulletin 1972;77:361–372.

Kawakami K, Dovidio JF, Moll J, Hermsen S, Russin A. Just say no (to stereotyping): Effects of training on the negation of stereotypic associations on stereotype activation. Journal of Personality and Social Psychology 2000;78:871–888. [PubMed: 10821195]

Klauer, KC.; Musch, J. Affective priming: Findings and theories. In: Klauer, KC.; Musch, J., editors. The psychology of evaluation: Affective processes in cognition and emotion. Mahwah, NJ: Erlbaum; 2003. p. 7-50.

Kruglanski, AW. Lay epistemics and human knowledge. New York: Plenum; 1989.

LaPiere RT. Attitudes vs. actions. Social Forces 1934;13:230–237.

Lord, CG.; Lepper, MR. Attitude Representation Theory. In: Zanna, MP., editor. Advances in experimental social psychology. Vol. 31. San Diego: Academic Press; 1999. p. 265-343.

Lowery BS, Hardin CD, Sinclair S. Social influence effects on automatic racial prejudice. Journal of Personality and Social Psychology 2001;81:842–855. [PubMed: 11708561]

MacCorquodale K, Meehl E. On a distinction between hypothetical constructs and intervening variables. Psychological Review 1948;55:95–107. [PubMed: 18910284]

Moors A, De Houwer J. Automaticity: A Theoretical and Conceptual Analysis. Psychological Bulletin 2006;132:297–326. [PubMed: 16536645]

Newby-Clark IR, McGregor I, Zanna MP. Thinking and caring about cognitive inconsistency: When and for whom does attitudinal ambivalence feel uncomfortable? Journal of Personality and Social Psychology 2002;82:157–166. [PubMed: 11831406]

Olson MA, Fazio RH. Implicit attitude formation through classical conditioning. Psychological Science 2001;12:413–417. [PubMed: 11554676]

Olson MA, Fazio RH. Implicit acquisition and manifestation of classically conditioned attitudes. Social Cognition 2002;20:89–103.

Olson MA, Fazio RH. Reducing the influence of extra-personal associations on the Implicit Association Test: Personalizing the IAT. Journal of Personality and Social Psychology 2004;86:653–667. [PubMed: 15161392]

Olson MA, Fazio RH. Reducing automatically-activated racial prejudice through implicit evaluative conditioning. Personality and Social Psychology Bulletin 2006;32:421–433. [PubMed: 16513796]

Petty, RE.; Cacioppo, JT. The elaboration likelihood model of persuasion. In: Berkowitz, L., editor. Advances in experimental social psychology. Vol. 19. San Diego, CA: Academic Press; 1986. p. 123-205.

Petty, RE.; Haugtvedt, C.; Smith, SM. Elaboration as a determinant of attitude strength: Creating attitudes that are persistent, resistant, and predictive of behavior. In: Petty, RE.; Krosnick, JA., editors. Attitude strength: Antecedents and consequences. Mahwah, NJ: Erlbaum; 1995. p. 93-130.

Powell MC, Fazio RH. Attitude accessibility as a function of repeated attitudinal expression. Personality and Social Psychology Bulletin 1984;10:139–148.

Reber R, Winkielman P, Schwarz N. Effects of perceptual fluency on affective judgments. Psychological Science 1998;9:45–48.

Richeson JA, Ambady N. Effects of situational power on automatic racial prejudice. Journal of Experimental Social Psychology 2003;39:177–183.

Roskos-Ewoldsen DR, Fazio RH. On the orienting value of attitudes: Attitude accessibility as a determinant of an object's attraction of visual attention. Journal of Personality and Social Psychology 1992;63:198–211. [PubMed: 1403611]

Rudman LA, Ashmore RD, Gary ML. "Unlearning" automatic biases: The malleability of implicit prejudice stereotypes. Journal of Personality and Social Psychology 2001;81:856–868. [PubMed: 11708562]

Rudman LA, Lee MR. Implicit and explicit consequences of exposure to violent and misogynous rap music. Group Processes & Intergroup Relations 2002;5:133–150.

Salancik GR, Conway M. Attitude inferences from salient and relevant cognitive content about behavior. Journal of Personality and Social Psychology 1975;32:829–840.

Sanbonmatsu DM, Fazio RH. The role of attitudes in memory-based decision making. Journal of Personality and Social Psychology 1990;59:614–622. [PubMed: 2254847]

Sanbonmatsu DM, Posavac SS, Vanous S, Ho EA, Fazio RH. The deautomatization of accessible attitudes. Journal of Experimental Social Psychology 2007;43:365–378.

Schuman, J.; Presser, S. Questions and answers in attitude surveys: Experiments in question form, wording, and context. New York: Academic Press; 1981.

Schwarz, N.; Bohner, G. The construction of attitudes. In: Tesser, A.; Schwarz, N., editors. Blackwell handbook of social psychology: Intraindividual processes. Malden, MA: Blackwell; 2001. p. 436-457.

Seibt B, Hafner M, Deutsch R. Prepared to eat: How immediate affective and motivational responses to food cues are influenced by food deprivation. European Journal of Social Psychology 2007;37:359–379.

Sherman SJ, Rose JS, Koch K, Presson CC, Chassin L. Implicit and explicit attitudes toward cigarette smoking: The effects of context and motivation. Journal of Social & Clinical Psychology 2003;22:13–39.

Shiffrin, RM.; Dumais, ST. The development of automatism. In: Anderson, JR., editor. Cognitive skills and their acquisition. Hillsdale, NJ: Erlbaum; 1981. p. 111-140.

Shook NJ, Fazio RH, Eiser JR. Attitude generalization: Similarity, valence, and extremity. Journal of Experimental Social Psychology. in press

Sinclair L, Kunda Z. Reactions to a Black professional: Motivated inhibition and activation of conflicting stereotypes. Journal of Personality and Social Psychology 1999;77:885–904. [PubMed: 10573871]

Smith ER, Fazio RH, Cejka MA. Accessible attitudes influence categorization of multiply categorizable objects. Journal of Personality and Social Psychology 1996;71:888–898. [PubMed: 8939039]

Strack F, Martin LL, Stepper S. Inhibiting and facilitating conditions of the human smile: A nonobtrusive test of the facial feedback hypothesis. Journal of Personality and Social Psychology 1988;54:768–777. [PubMed: 3379579]

Tetlock PE, Kim I. Accountability and judgment processes in a personality prediction task. Journal of Personality and Social Psychology 1987;52:700–709. [PubMed: 3572733]

Thurstone LL. Attitudes can be measured. American Journal of Sociology 1928;33:529–554.

Tourangeau R, Rasinski A. Cognitive processes underlying context effects in attitude measurement. Psychological Bulletin 1988;103:299–314.

Uleman, JS. Spontaneous versus intentional inferences in impression formation. In: Chaiken, S.; Trope, Y., editors. Dual-process theories in social psychology. New York: Guilford; 1999. p. 141-160.

Wallace DS, Paulson RM, Lord CG, Bond CF. Which Behaviors Do Attitudes Predict? Meta-Analyzing the Effects of Social Pressure and Perceived Difficulty. Review of General Psychology 2005;9:214–227.

Wicker AW. Attitudes versus actions: The relationship of verbal and overt behavioral responses to attitude objects. Journal of Social Issues 1969;25:41–78.

Wilson TD, Hodges SD, LaFleur SJ. Effects of introspecting about reasons: Inferring attitudes from accessible thoughts. Journal of Personality and Social Psychology 1995;69:16–28. [PubMed: 7643298]

Wilson TD, Lindsey S, Schooler TY. A model of dual attitudes. Psychological Review 2000;107:101–126. [PubMed: 10687404]

Wittenbrink B, Judd CM, Park B. Spontaneous prejudice in context: Variability in automatically activated attitudes. Journal of Personality and Social Psychology 2001;81:815–827. [PubMed: 11708559]

Zaller, J.; Feldman, S. American Journal of Political Science. Vol. 36. Hillsdale, NJ: Erlbaum; 1992. A simple theory of the survey response: Answering questions versus revealing preferences; p. 579-616.

Zanna, MP.; Fazio, RH. The attitude-behavior relation: Moving toward a third generation of research. In: Zanna, MP.; Higgins, ET.; Herman, CP., editors. Consistency in social behavior: The Ontario symposium. Vol. 2. Hillsdale, N.J.: Erlbaum; 1982. p. 283-301.

Zanna, MP.; Fazio, RH.; Ross, M. The persistence of persuasion. In: Schank, RC.; Langer, E., editors. Beliefs, reasoning, and decision making: Psycho-logic in honor of Bob Abelson. Hillsdale, NJ: Erlbaum; 1994. p. 347-362.

Zanna, MP.; Rempel, JK. Attitudes: A new look at an old concept. In: Bar-Tal, D.; Kruglanski, AW., editors. The Social Psychology of Knowledge. Cambridge, England: Cambridge University Press; 1988. p. 315-334.