

THE META-COGNITIVE MODEL (MCM) OF ATTITUDES: IMPLICATIONS FOR ATTITUDE MEASUREMENT, CHANGE, AND STRENGTH

Richard E. Petty
Ohio State University

Pablo Briñol
Universidad Autónoma de Madrid

Kenneth G. DeMarree
Ohio State University

We present a Meta-Cognitive Model (MCM) of attitudes. According to the MCM, an attitude object can be associated with both positive and negative evaluations that can be detected with modern implicit (automatic) measures of attitudes. These evaluative associations can be further associated with validity tags that are also consulted when completing deliberative attitude measures. We explain how the MCM accounts for existing findings in the attitudes literature and discuss recent studies and novel predictions derived from the MCM framework. Particular attention is devoted to factors that produce discrepancies between automatic and deliberative measures of attitudes and their consequences. Implications for attitude change, attitude strength, and the domain of prejudice are also considered.

Attitudes are the bedrock of social psychology and a pervasive concept throughout the social sciences. How, then, should they be conceptualized? Over the many years of its existence, the attitude concept has had many different definitions (see Eagly & Chaiken, this issue). Attitudes have been described as hypothetical constructs at times, and at other times as real (e.g., see Krosnick, Judd, & Wittenbrink, 2005). They have been assumed to be conscious and unconscious (e.g., see Greenwald & Banaji, 1995). They have referred to behavioral, cognitive, and/or emotional reactions (e.g., see Zanna & Rempel, 1988). Despite these variations,

Please address correspondence to Richard E. Petty, Department of Psychology, The Ohio State University, 1835 Neil Avenue, Columbus, OH 43210-1222; E-mail: petty.1@osu.edu.

one common thread across the years is that attitudes were assumed to have an evaluative component. Indeed, today attitudes commonly refer to people's evaluations of a wide variety of objects, issues, and people, including the self. Our approach is compatible with this view.

Our primary goal in this article is to present our Meta-Cognitive Model (MCM) of attitudes (see also Petty, 2006; Petty & Briñol, 2006). Before describing the MCM in some detail, it is useful to outline briefly three alternative views of attitudes—one that is now classic and two others that are of more recent vintage. After describing these alternatives and the MCM, we compare our approach to the alternatives. Then, we articulate what the MCM can contribute to understanding what deliberative and automatic attitude measures reflect, what it means when these measures conflict, what happens when attitudes change, how the MCM relates to the underlying strength of attitudes, and some other unique perspectives offered by the MCM.¹

THREE PERSPECTIVES ON ATTITUDE REPRESENTATION

There is one model of attitude representation that has dominated the literature for some time, and two that have gained adherents more recently. We describe each approach briefly so that we can compare them to our meta-cognitive framework.

SINGLE ATTITUDE MODEL

The first, and now classic approach to attitude representation, asserts that attitudes are best conceptualized as an object-evaluation link in

1. We use the terms deliberative and automatic attitudes and attitude measures rather than implicit/explicit attitudes/measures because of the ambiguity that surrounds the use of the latter terms in the literature (Petty, Tormala, Briñol, & Jarvis, 2006). In general, researchers use the term "implicit attitude" to refer to what is assessed by measures that tap into relatively quick (even automatic) evaluative associations (i.e., implicit measures), whereas the term "explicit attitude" refers to what is assessed with more deliberative and undisguised measures (explicit measures). Sometimes, however, researchers use the term "implicit" to refer to an attitude (or measure) of which people are unaware, or to a basis of the attitude of which people are unaware, or to an effect of an attitude of which people are unaware (see also Fazio & Olson, 2003; Petty, Wheeler, & Tormala, 2003). To allow for more precision, we will use the terms "automatic" and "deliberative attitudes" that are tapped by automatic and deliberative measures (though neither measure is process pure, e.g., see Sherman, in press). In earlier literature, the term "opinion" was often used to refer to the verbal expression of an implicit (internal) attitude (e.g., see Hovland, Janis, & Kelley, 1953; Thurstone, 1928).

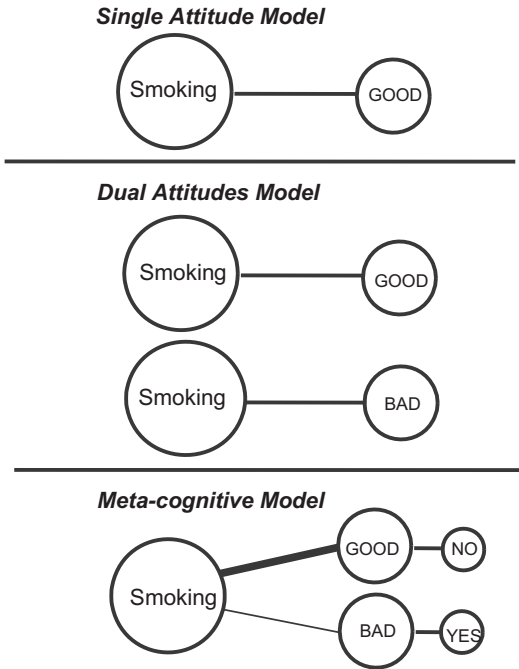


FIGURE 1. Structural Representation of Three Attitude Models.

memory (e.g., Fazio, 1995; Fiske & Pavelchak, 1986; see also Fazio, this issue). This approach, which we call the single attitude model, is depicted in the top panel of Figure 1 and shows a person who has a positive attitude toward smoking. The most well-articulated and influential example of this conceptualization is Fazio’s MODE Model (Fazio & Towles-Schwen, 1999), which has contributed much to our understanding of attitude structure and expression. In brief, the MODE Model holds that people have stored evaluative associations to attitude objects. In this framework, automatic measures of attitudes (e.g., evaluative priming; Fazio, Jackson, Dunton & Williams, 1995; and the Implicit Association Test [IAT]; Greenwald, McGhee & Schwarz, 1998) tend to assess the stored evaluative association, whereas more deliberative measures (e.g., semantic differential; Osgood, Suci, & Tannenbaum, 1957) tap the retrieved evaluative association along with the outcome of any downstream cognitive processes. Thus, if a person expresses a different attitude on a deliberative than an automatic measure, it is presumably because he or she has engaged in some thought that modifies the initial automatic evaluative

reaction that comes to mind (see also, Gawronski & Bodenhausen, 2006, for similar assumptions). This thought can reflect additional mental contents that are brought to mind or activated by the context, or it can stem from impression management or correction motives (see Fazio, this issue).

DUAL ATTITUDES MODEL

A second approach that has captured the attention of social psychologists more recently argues that people can hold separate explicit (conscious, deliberative) and implicit (unconscious, automatic) attitudes (e.g., Greenwald & Banaji, 1995; Wilson, Lindsey, & Schooler, 2000) which can take on different values. Although there are various assumptions that have been made about dual attitudes by the different theorists who advocate this view, three assumptions are fairly common.

Perhaps the most important assumption shared by many dual attitudes proponents is that the dual attitudes (implicit and explicit) have separate mental representations (e.g., see Wilson et al., 2000). As depicted in the middle panel of Figure 1, implicit and explicit attitudes are viewed as distinct mental entities that are stored separately in different areas of the brain (e.g., see DeCoster, Banner, Smith, & Semin, 2006). Thus, a person might have a deliberative (explicit) attitude toward an object of one valence but an automatic (implicit) attitude of a different valence.²

A second common assumption is that the two attitudes stem from distinct mental processes. Implicit attitudes are said to stem from associative processes such as evaluative conditioning, whereas explicit attitudes stem from propositional processes such as thinking about message arguments (e.g., Rydell, McConnell, Mackie, & Strain, 2006). The different mental processes responsible for explicit and implicit attitudes are often characterized as being governed by separate mental systems (e.g., reflective/impulsive, fast/slow; DeCoster et al., 2006; Rydell & McConnell, 2006).

Third, implicit and explicit attitudes are postulated to be relatively independent and to operate in different situations. Thus, these attitudes are not expected to be in conflict (DeCoster et al., 2006), but instead are

2. In the single attitude approach, deliberative and automatic *measures* can show different attitudes, but it is not assumed that there are separate mental representations. Rather, a difference in the measures can occur when a deliberative attitude is "constructed" online from consideration of the automatic evaluation along with other things that come to mind (Fazio & Towles-Schwen, 1999).

postulated to work in different arenas. In particular, implicit attitudes are postulated to guide behavior in spontaneous situations when people are not engaged in much thought, whereas explicit attitudes are said to guide behavior when people are being reflective (see Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997, Dovidio, Kawakami, Smoak, & Gaertner, in press). In this sense, then, the manner of measuring implicit and explicit attitudes (i.e., with automatic versus deliberative measures) matches the situations in which they guide behavior (Vargas, 2004).

When considering all of these assumptions together, the dual attitudes framework suggests that attitudes assessed with automatic and deliberative measures are quite different. These attitudes have separate representations, are the result of different processes from separate mental systems, and operate in different situations. There is not much, if any, interaction between them. Indeed, they have been characterized as "two ships passing in the night" (Cohen & Reed, 2006, p. 9).

PROCESS MODELS OF ATTITUDES

One additional perspective that takes a radically different approach to attitudes holds that there are no stored evaluations in memory. Instead, according to this *constructivist* perspective, attitudes are formed, as needed, based on currently salient beliefs, feelings, and behaviors (e.g., Schwarz & Bohner, 2001; Wilson & Hodges, 1992; see also Schwarz, this issue). According to this approach, different contexts will make different knowledge accessible, resulting in changes in people's reported evaluations. Any consistency in attitudinal reports, according to this perspective, comes from the same set of building blocks being retrieved each time and being used in computing the current evaluation. In our view, it is not clear why a mental system would develop such that it would link particular attributes to evaluations (e.g., being lazy is bad) but not link the overall attitude object to an evaluation (e.g., Fred is bad; see Petty et al., 2003). Indeed, research suggests that people often do not retrieve attribute information when they have a previously formed and relevant attitude to guide decisions (see Lingle & Ostrom, 1981; Lichtenstein & Srull, 1985).

One recent instantiation of a process approach to attitudes is the APE (Associative Propositional Evaluation) model (Gawronski & Bodenhausen, 2006). This framework holds that people can respond positively or negatively to some attitude object based solely on the affect that is associated with the object, or based on the propositions that come to mind with respect to the object. Affect associated with an object can be translated into propositional form (e.g., I like this) and then checked for

validity by an online process that examines whether the evaluative proposition is consistent with other salient propositions. In this framework, there are no stored evaluations (attitudes) *per se*, only stored affect and beliefs (propositions) that serve as input to an expressed evaluation (see Gawronski & Bodenhausen, this issue).

THE META-COGNITIVE MODEL (MCM)

Our Meta-Cognitive Model (MCM) shares some features with each of the approaches just described, but also has notable differences. In brief, the MCM holds that attitude objects can be linked in memory to both positive and negative evaluations that can vary in the degree to which they are endorsed or not (see bottom panel of Figure 1). The MCM is based on a number of assumptions that we outline next.

First, in accord with the single attitude approach and in contrast to the constructivist view, the MCM holds that attitude objects can be linked in memory to global evaluative associations, and these associations can vary in strength (e.g., Fazio & Towles-Schwen, 1999; see also Fazio, this issue). Also parallel to the single attitude approach, features of the context can determine whether an activated evaluation is used. Thus, just as one can store one's birthdate in long-term memory and retrieve it when needed to construct how old one is currently (i.e., by considering one's birthdate along with today's date), people can retrieve evaluations associated with attitude objects and modify them according to the current situation (e.g., "I generally like cake, but this one smells odd, so I won't have any").

Second, in concert with the dual attitudes approach and recent research and theory on the separability of positivity and negativity (e.g., Cacioppo, Gardner, & Berntson, 1997), the MCM holds that attitude objects can sometimes be linked in memory to evaluative associations of opposite valence (see also, deLiver, van der Pligt, & Wigboldus, 2007). Whether a positive or negative evaluation comes to mind first will depend on all of the various factors that can affect memory. Thus, the number of prior positive and negative experiences, the recency of those experiences, and the context in which those experiences took place will matter. For example, if negative experiences with some object were formed largely at work but positive associations were formed at home, then association measures given in these different contexts should reveal different responses. Similarly, if people have experienced positive reactions to African Americans in a sports context, but negative reactions in an urban setting, measures of association that include these contextual features should show different evaluations (Barden, Maddux,

Petty, & Brewer, 2004; Wittenbrink, Judd, & Park, 2001). It is important to note that not all attitude objects are expected to be associated with both positive and negative associations. Indeed, when an attitude object is linked to only one evaluative association, an attitude structure similar to the single-attitude model emerges.

Third, in contrast with a common assumption of dual attitudes models, we do not assume that opposite evaluative associations, if present, necessarily stem from different underlying processes—that is, both positive and negative evaluations can stem from associative or from propositional processes. Thus, although it is possible to have one valenced association that stems from propositional processes while the other stems from associative processes as specified by some dual attitudes theorists, it is also possible for both valenced associations to be based on simple associative processes or more reflective thought processes or some combination of each.

Fourth, the feature of the MCM that gives the model its name is the assumption that people can tag their evaluative associations as true or false, or that they are held with varying degrees of confidence. In this way, the model builds on existing research on meta-cognition (Jost, Kruglanski, & Nelson, 1998). Meta-cognition refers to thoughts about thoughts or thought processes. Research in this domain has demonstrated that a person can have an initial or primary thought (e.g., “I like candy”) that is further reflected upon by a meta-cognitive or secondary thought (e.g., “I am confident that I like candy”; see Petty, Briñol, Tormala, & Wegener, 2007). The meta-cognitive associations in the MCM can be represented in various ways such as yes/no, confidence/doubt, true/false, accept/reject, and so forth.³ Furthermore, these meta-cognitions can vary in the strength of their association to the linked evaluation, and the strength of this association will determine the likelihood that the perceived validity of an evaluation will be retrieved along with the evaluation itself.

Thus, in concert with the APE model (Gawronski & Bodenhausen, 2006) and other conceptualizations (e.g., Cohen & Reed, 2006; Kruglanski, 1989; Petty, Briñol, & Tormala, 2002), the MCM puts a focus on validity processes. Unlike the APE model, however, we do not assume that validity stems from consistency analyses alone, but rather that

3. Affective validation is also possible wherein people's attitudes make them happy or sad, comforted or anxious. Although we focus on validity tags in this model, we also acknowledge that other tags might exist, and these tags might also exert an impact on attitudinal processes. For example, a person might tag a negative racial evaluation as “inappropriate to express” even though he or she might personally endorse the association.

confidence can be inferred directly (e.g., from ease of attitude retrieval; see Holland, Verplanken, & van Knippenberg, 2003), or result from various inferences based on the perceived content of one's attitude (e.g., negatively framed attitudes are held with more confidence than positive ones; Bizer & Petty, 2005) or the process by which an attitude was formed, maintained, or changed (e.g., having difficulty defending one's attitude can lead to less confidence; Petty, Tormala, & Rucker, 2004), among other factors. Most notably, however, the MCM goes beyond the idea that validation is solely an online process and holds that perceived validities, like the evaluations themselves, can be stored for later retrieval. In other words, the MCM assumes that just as it is adaptive to store evaluations to guide decision making and action (Fazio, 1995), so too is it adaptive to know if any activated evaluation is a reliable guide. The storage of validity information can presumably make decision making and action more efficient. To the extent that the retrieval of validity tags becomes automatic, it even becomes possible for people to quickly correct for undesired evaluations that might come to mind. Furthermore, just as evaluative associations can be context specific, so too can people learn to associate invalidity tags with evaluations more in some contexts than others (e.g., see Maddux, Barden, Brewer, & Petty, 2005).

Although there is no definitive research on the storage of validity tags, some evidence consistent with this idea comes from research on the stability of belief certainty over time. For example, in one study, a manipulation of expressed agreement with one's judgment by another person (i.e., social consensus) produced the same increase in judgmental confidence whether that confidence was measured immediately or 48 hours later (Wells, Olson, & Charman, 2003). In other research, attitude certainty measured at one point in time was shown to predict certainty-related outcomes (i.e., resistance of attitudes to change) at a later time, even when certainty was not made salient at the second occasion (e.g., Bassili, 1996, study 2). These results are consistent with the idea that confidence can reside in memory. The MCM does not specify exactly how evaluations and validities are stored in memory, however. In other words, memory for evaluations and their validities can be conceptualized as part of a traditional semantic association network (e.g., Fiske & Pavelchak, 1986) or as linked patterns of activation in a connectionist model (e.g., Eiser, Fazio, Stafford, & Prescott, 2003; see also Conrey & Smith, this issue). Either framework can accommodate the postulated linkages in the MCM (e.g., see Van Overwalle & Siebler, 2005 for a connectionist model wherein attitude objects are linked to both positivity and negativity).

Finally, the MCM concurs with research on cognitive negation that suggests that untagged evaluations are presumed to be true unless evidence

against them is or has been generated (e.g., see Gilbert, 1991). Furthermore, research on negation suggests that successful negation is quite difficult (e.g., Deutsch, Gawronski, & Strack, 2006). Indeed, overriding one's negated attitudes will require motivation and ability, at least in the early stages (Betsch, Haberstroh, Molter, & Glöckner, 2004). People could attempt to invalidate their previously formed evaluative associations for many reasons. For example, people might reject an evaluative association because they realize that it stems from the culture (e.g., media exposure) and not from personal beliefs (e.g., Devine, 1989). In addition, the association can represent the opinions of others that have been encoded (e.g., Han, Olson, & Fazio, 2006; Priester & Petty, 2001). Also, the association can represent a previously accepted personal view that has more recently been discredited (e.g., Gregg, Seibt, & Banaji, 2006; Petty et al., 2006).⁴ When the association and negation are presented at the same point in time (e.g., John is not clean), people can sometimes reverse the association (i.e., "not clean" becomes "messy;" see Mayo, Schul, & Burnstein, 2004), but when the negation follows the association in time (e.g., John is clean. . . . False), this is less likely.

We believe that the MCM provides a more complete, integrative, and flexible view of a person's underlying attitude structure than alternatives, and provides an explanation for various attitudinal phenomena such as how one underlying attitude structure can lead to different outcomes on deliberative versus automatic measures. The MCM also has some unique implications. We discuss these issues next.

UNDERSTANDING AUTOMATIC VERSUS DELIBERATIVE ATTITUDE ASSESSMENTS

Much attention has been paid recently to discrepancies that can emerge between attitudes assessed with deliberative versus automatic measures. These differences have emerged both when an attempt is simply made to assess a person's existing attitudes (e.g., the person scores positive toward minorities on a deliberative measure but negative on an automatic measure; see Greenwald et al., 1998), and when one attempts to assess the effectiveness of an attitude change manipulation (e.g., the person shows attitude change from the perspective of one type of measure but not the other; see Gawronski & Bodenhausen, 2006). We discuss each of these kinds of discrepancies in turn.

4. When the negated evaluation is a prior attitude, we have referred to our approach as the PAST (Prior Attitudes are Still There) Model (Petty & Jarvis, 1998; Petty et al., 2006). That is, the PAST Model is a specialized case of the more general MCM.

DIVERGENCE IN AUTOMATIC VERSUS DELIBERATIVE ATTITUDE MEASURES

Imagine a situation in which a person's attitude shows a positive evaluation of some attitude object on a measure of automatic evaluation, but a negative evaluation on a deliberative measure. According to the single attitude model, the underlying structure is best captured by the top panel of Figure 1—that is, at the association level, the attitude object is linked to a positive evaluation and this shows up on the automatic measure. If a negative evaluation is expressed on a deliberative measure, according to this approach, it must be due to downstream cognitive processes. For example, the person could bring to mind various negative attributes of the object that override the positive association, or it could be that a person's egalitarian motives or desires to be unprejudiced lead to a correction of the negative underlying attitude (e.g., Dunton & Fazio, 1997; Plant & Devine, 1998). According to the dual attitudes perspective, it is possible that the person has distinct positive and negative evaluations residing in separate memory systems (Figure 1, middle panel) that are called forth on the different types of measures. The constructivist perspective does not allow for stored general evaluations and thus holds that different material from memory is retrieved and integrated on automatic versus deliberative measures.

Although the MCM accepts that the processes outlined by the single and constructivist positions can sometimes occur (i.e., when only one evaluative association, or none, is stored), it also points to another possibility—that is, one possible attitude structure leading to divergence in automatic versus deliberative measures is represented in the bottom panel of Figure 1. The MCM assumes that contemporary measures of automatic evaluation tap (although not perfectly) into evaluative associations largely without respect to validity tags. Thus, if such a measure reveals an overall positive attitude, the MCM holds that this is because positive associations are stronger than negative ones without respect to the validity tags that might be attached to those evaluations. The reverse is the case if such measures reveal an overall negative attitude. Explicit measures, in comparison, reflect evaluative associations as modified by stored validity tags as well as other considerations that come to mind prior to responding.

So, in the evaluative structure depicted in the bottom panel of Figure 1, which attitude would be assessed with each measure? An automatic measure would reveal a somewhat positive attitude toward the object. This is because although the object is associated with both good and bad, the strength of association to good is stronger than to bad. It does not matter

that the positive evaluation is rejected because such negations are generally not tapped by measures of automatic evaluative association. In contrast, because deliberative measures consider stored (as well as online) validity, such measures are more likely to show a negative attitude.

There are several reasons that the impact of validity tags will be difficult to observe on automatic measures of attitudes. First, because a validity tag is a stored form of meta-cognition (i.e., secondary cognition), it is not directly linked to the attitude object, but is instead linked to the evaluative association (i.e., the primary cognition), which is in turn linked to the attitude object. Because of this, validity tags will take more time to retrieve than evaluations, and the impact of these associations are less likely to be evident on automatic attitude measures. In addition, there are many circumstances where validity tags will not be as strongly linked to the evaluation as the evaluation is to the attitude object (e.g., because less thought was devoted to forming the validity association than the evaluation). However, as the strength of the link between an evaluation and the associated validity tag increases, the likelihood that it will be retrieved increases.

ATTITUDE CHANGE

Another phenomenon that has been observed is that when people are subjected to an attitude change manipulation, sometimes attitudes appear to change more on deliberative than automatic measures (e.g., Gregg et al., 2006), sometimes they appear to change more on automatic than deliberative measures (e.g., Karpinski & Hilton, 2001), sometimes the measures are affected differently by different aspects of the persuasion treatment (DeCoster et al., 2006; Rydell & McConnell, 2006), and sometimes the measures are influenced in a similar manner by the manipulation (see Gawronski & Bodenhausen, 2006, for a comprehensive review). How does the MCM account for each of these situations? First, we note that the simplest situation is one in which both automatic and deliberative measures show change to some persuasion treatment. This is because both associative (low-effort) and propositional (high-effort) processes should be capable of forging evaluative associations (Briñol, Petty, & McCaslin, in press). To the extent that this occurs, both automatic and deliberative measures would show change. However, according to the MCM, it is also possible for some persuasion treatment to affect one type of measure but not the other, or affect each differently. We explain these situations next.

More Change on Deliberative than Automatic Measures. Consider first a situation in which deliberative attitudes will likely show greater change

than automatic ones. Imagine a person who has favored smoking for a long time and whose attitude can be conceptualized as in the top panel of Figure 2. This person begins with both deliberative and automatic attitudes in synch—both are positive. Then, this person begins to receive numerous anti-smoking messages that are convincing. The person now develops an ambivalent attitude in that both positive and negative associations are present, and endorsed, as depicted in the second panel of Figure 2. When measured with a bipolar deliberative measure (good–bad), the person would show a neutral or slightly positive attitude that represents the integration of the separate positive and negative reactions. However, if desired, the separate positive and negative reactions could be assessed with deliberative unipolar measures that asked about positivity and negativity separately (see Kaplan, 1972). Like the bipolar deliberative measure, in the realm of automatic attitude assessments, the ambivalent person would come out neutral or slightly positive on a measure that collapsed across positivity and negativity (e.g., IAT; Greenwald et al., 1998). However, if positivity and negativity were assessed separately, one could see that for the ambivalent person, there are automatic associations to both positivity and negativity (e.g., de Liver et al., 2007; Newby–Clark, McGregor, & Zanna, 2002)


So, at the point of receiving and accepting negative information about smoking, the automatic and deliberative measures are again showing comparable effects. Over time, however, the person may come to completely reject his or her initial positive evaluation of smoking and possess the attitude structure shown in Panel 3 of Figure 2. Once the person rejects the initial attitude, at this point the deliberative measure will likely show more change from its initial value than the automatic measure. Because the automatic measure does not consider the validity of the associations (unless the negation itself becomes automatic), it will show the same pattern of mixed reactions as before, whereas the explicit measure will now indicate a negative reaction toward smoking.

Evidence for a pattern of greater change in a deliberative than an automatic measure was obtained in a study by Petty et al. (2006, study 1). In this study, participants first formed an initial positive or negative attitude toward another person via evaluative conditioning. This manipulation was effective in modifying both automatic (evaluative priming) and deliberative (semantic differential) measures of attitudes. Next, participants received information about the opinions of the target person on several important issues that would make the person appear either very likable (i.e., had similar attitudes to the participant) or dislikable (i.e., had dissimilar attitudes to the participant; see Byrne, 1961). In some conditions, this information reinforced the initial impression (i.e., no attitude change) and in other conditions this information contradicted the


1. Univalence



2. Explicit Ambivalence



3. Implicit Ambivalence (a)



4. Implicit Ambivalence (b)




FIGURE 2. Depiction of Univalence, Explicit Ambivalence, and Implicit Ambivalence from the Perspective of the Meta-Cognitive Model.

initial impression (i.e., attitude change). In the reinforcement conditions, both the deliberative and automatic measures showed the same pattern of results (i.e., more positive attitudes toward the similar person who was conditioned positively than to the dissimilar person who was conditioned negatively). However, in the contradiction (attitude change) conditions, the measures diverged such that attitudes were more sensitive to the contradictory similarity information about the target on the deliberative than on the automatic measure. In this research, the deliberative measure reflected the fact that the old attitude was rejected, whereas the automatic measure reflected fast association of the target person to both the old and the new evaluations. This state of affairs represents what might be called the normal attitude change situation in which people reject their previous attitude and accept a new one. The explicit measure tracks this change quite well, but the implicit measure lags behind

because of its relative insensitivity to the negation (see also Gregg et al., 2006).

More Change on Automatic than Deliberative Measures. Notably, according to the MCM, in some attitude change situations it is possible for an implicit (automatic) attitude measure to tap change better than an explicit (deliberative) measure. Imagine our smoker again (Panel 1 of Figure 2). This time the smoker is exposed to an advertising campaign that features disgusting photos of black lungs, cancerous growths, and other negative consequences of smoking. The smoker starts to notice an automatic negative feeling whenever opening a cigarette pack. Then, a friend points out the manipulative advertising campaign and the smoker recognizes that this is likely the cause of the negative feeling. As a consequence, he rejects the campaign and the attitude it implies. In fact, the smoker resents being targeted by the messages. If an explicit attitude measure toward smoking is administered at this point, it will likely show little change—the person still reports being quite positive toward smoking. Yet, an implicit measure would likely show a less positive reaction than previously due to the new negative associations. The structure of this person's attitude according to the MCM is presented in Panel 4 of Figure 2.

This type of attitude change situation is a bit more unusual than the previous one, but the literature on attitude change has suggested a number of situations in which it can arise. Perhaps the most obvious situation that contains this possibility is work on the sleeper effect (see Kumkale & Albaracín, 2004 for a review). In this paradigm, people are exposed to a very strong persuasive message (counter to a person's initial attitude), but this strong message is then discounted (e.g., claimed to be false or attributed to a low credible source). The presence of a new evaluative association (from the strong message) should produce change in the automatic measure, but the negation of the new association (from the discounting cue) should leave the deliberative measure unchanged. Note that in this situation, it is the automatic measure that taps the new (more recent) attitude and the deliberative measure that taps the old (initial) attitude. Thus, the MCM does not predict that automatic measures invariably tap "old" attitudes, whereas deliberative measures tap more recent ones. Rather, it is the strength of the evaluative association that matters along with validity tags. When the normal attitude change paradigm is modified so that a strong new attitude is first created and then negated, the automatic measure reflects the more recent attitude, whereas the deliberative measure reflects the older attitude (since the new one is negated).

In this situation, although the deliberative measure does not show change initially, over time change can emerge (i.e., the sleeper effect). Viewed from the perspective of the MCM, what happens over time is

that the negation link is weakened (forgotten or dissociated) so that the evaluative association formed at the time of message exposure can have an impact on the deliberative measure. If the negation link is strongly attached to the evaluative association (e.g., because this link was activated frequently or with a great deal of thought), however, the sleeper effect may not emerge. Indeed, although it has never been tested, the MCM anticipates that in a sleeper effect paradigm, those who show the largest discrepancy between automatic and deliberative measures right after the message (i.e., large change on automatic and small change on deliberative) would show the largest sleeper effect. Another paradigm in which people may consciously negate an evaluative association that has formed (allowing for delayed influence on the deliberative measure) is that of minority influence (see Tormala, DeSensi, & Petty, 2007).

Different Impact on Automatic and Deliberative Measures. As our final situation, we turn to research in which one manipulation affects automatic and deliberative measures differently. Consider a study by DeCoster et al. (2006) in which participants were presented with photographs of various target individuals accompanied by an evaluative assertion regarding the target (e.g., Sam is smart) or a negation (e.g., Jack is *not* smart). Coming from a dual attitudes perspective, the authors argued that the fast learning (propositional) system would be capable of handling negations and thus would be able to differentiate the two individuals. The slow learning (associative) system, on the other hand, would not be able to process negations and thus would associate both Sam and Jack with intelligence. With respect to automatic versus deliberative measures, the dual systems approach predicts that when presented with a negation, an automatic measure would assess Jack as intelligent, but a deliberative measure would assess him as not smart (or as stupid). Because the automatic and deliberative measures used in their research indicated opposite evaluative associations to Jack, the authors reasoned that negations cause people to form dual attitudes that are stored in separate memory systems.

The MCM explains these data differently—that is, rather than having two representations of Jack as both smart and stupid stored in separate systems as depicted in the top panel of Figure 3, the MCM holds that the information is encoded as depicted in the bottom panel of Figure 3. Therefore, Jack is associated with intelligence, but this is negated. For the reasons already outlined, this cognitive structure would lead to the prediction that an automatic measure would assess Jack as intelligent because this is directly linked to Jack, but motivation and ability are needed to retrieve the negation. Thus, a deliberative measure would assess Jack as stupid because on the more thoughtful measure, it is more likely that the negation would be retrieved.

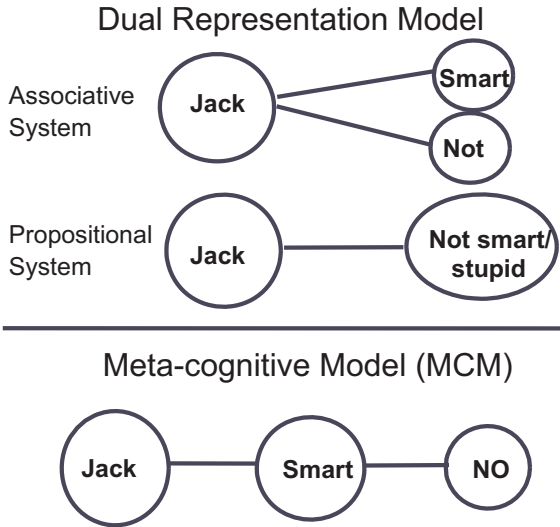


FIGURE 3. Representation of “Jack Is Not Smart” from the Dual Systems Approach and the MCM.

ADDITIONAL IMPLICATIONS OF THE META-COGNITIVE MODEL

It is important for the MCM to be able to account for phenomena that exist and that have been explained by other theories. If the MCM could not account for phenomena such as those just reviewed, it would not be very useful. However, the MCM goes beyond providing an alternative account for empirical findings that other theories can also explain. It also makes unique predictions and we turn to these next.

Earlier in this article we explained how the MCM could account for discrepancies between automatic and deliberative measures of attitudes as well as deliberative measures changing more to an influence attempt than automatic measures and vice versa. The MCM makes a unique prediction about both of these situations that we consider next. After this, we consider what the MCM says about the strength of deliberative attitude reports and their correlation with automatic measures, and we present some new insights that the MCM can offer research on prejudice.

AUTOMATIC–DELIBERATIVE DISCREPANCIES

The MCM notes that automatic–deliberative discrepancies can occur when people have both positive and negative associations to an attitude object, and one of these associations is rejected. Although people with this attitude structure (see the bottom panel of Figure 1) would experience no explicit ambivalence because they do not endorse both positive and negative aspects of the attitude object, there would still be ambivalence at the level of evaluative associations—an ambivalence that we have referred to as *implicit ambivalence* (Petty et al., 2006). This ambivalence is implicit in that it is at the level of automatic associations and the person does not experience it consciously as ambivalence—that is, there is no explicit conflict because one of the valences is explicitly rejected. If there is some implicit ambivalence, then people might be expected to attempt to resolve it because evaluative conflict (even if just at the associative level) is typically unpleasant.

As explained earlier, according to some versions of the dual attitudes approach, there should not be any conflict when automatic and deliberative measures diverge because the implicit and explicit attitudes are assumed to have separate lives. Either one or the other will guide responses depending on whether the situation is a deliberative or a spontaneous one (Dovidio et al., in press). The two evaluations should not be jointly activated in any given situation, which is a requirement for ambivalence (Newby–Clark et al., 2002). According to the dual attitudes approach, the two attitudes are the result of separate mental systems and as DeCoster and colleagues (2006) noted, “the fact that the two systems store their representations in separate areas of the brain means that any inconsistencies between them do not have to be resolved” (p. 9). However, in accord with the MCM, we have obtained evidence that people act as if automatic–deliberative divergence produces a state of (implicit) ambivalence.

In a series of studies we have shown that discrepancies between automatic and deliberative measures of attitudes can lead to some of the same consequences as explicit endorsement of positive and negative attributes of some attitude issue or object. One well-known consequence of explicit attitudinal ambivalence is that it leads to enhanced information processing in a presumed attempt to resolve the ambivalence. In one study, for example, Maio, Bell, and Esses (1996) measured participants’ explicit ambivalence regarding the issue of immigration to Canada (i.e., the extent to which they endorsed both positive and negative reactions to the issue), and then exposed them to a message favoring immigration from Hong Kong to Canada that contained either strong or

weak arguments. The degree to which participants processed the message information was assessed by examining the extent to which the quality of the arguments affected post-message attitudes toward immigration (Petty, Wells, & Brock, 1976). When people are thinking carefully about information, they should be affected by the quality of the arguments that a message contains (see Petty & Cacioppo, 1986). As hypothesized, Maio et al. (1996) found that individuals who had explicitly ambivalent attitudes toward immigration were more influenced by argument quality than were individuals low in ambivalence, suggesting that they engaged in enhanced scrutiny of the information.

In one study testing the notion that automatic–deliberative attitude discrepancy could lead to enhanced information processing (Briñol, Petty, & Wheeler, 2006), we assessed undergraduates' self-evaluation with both automatic (IAT; Greenwald et al., 1998) and deliberative (Rosenberg, 1965) measures and then calculated the absolute value of the difference between the two standardized measures as our index of discrepancy (see also Kehr, 2004). Next, participants were exposed to either a strong or weak message about eating vegetables that was framed as self-relevant (i.e., relevant to one's personal lives and thus relevant to the discrepancy) or not. The results are displayed in Figure 4. When the message was framed as self-relevant, the extent of automatic–deliberative discrepancy interacted with argument quality to affect attitudes. The greater the automatic–deliberative discrepancy, the more participants differentiated strong from weak arguments (top panel). However, when the same strong and weak messages were framed as irrelevant to the self (i.e., the message was said to be about the properties of vegetables), discrepancy did not interact with argument quality to predict attitudes (bottom panel). This suggests that automatic–deliberative discrepancies do not lead to motivation to process all information—only that relevant to the discrepancy.

Finally, we have also attempted to provide more direct evidence for the idea that automatic–deliberative discrepancies produce implicit ambivalence (Briñol, Petty, & Wheeler, 2003). Specifically, in one study we found that as discrepancy in automatic versus deliberative self-esteem (as assessed using the absolute value of the difference between participants' standardized automatic and deliberative self-esteem scores) increased, the strength with which participants automatically associated doubt-words with self-words on an IAT also increased. However, increased discrepancy was not associated with explicit reports of self-uncertainty. This suggests that the self-doubt that accompanies automatic–deliberative discrepancy is either not open to conscious awareness, or is explicitly denied or discounted.




FIGURE 4. Interaction of Frame, Implicit-Explicit Self-esteem Discrepancy, and Argument Quality on Attitudes (Data from Briñol, Petty, & Wheeler, 2006, Experiment 4).

ATTITUDE CHANGE

As noted in our earlier discussion of persuasion, the MCM holds that in situations of changing attitudes from one valence to another, there is the potential for implicit ambivalence. This is because when the validity tags are ignored, both positive and negative evaluations are associated with the attitude object (see Panels 3 and 4 in Figure 2). Is there any evidence that attitude change situations can produce implicit ambivalence in the absence of explicit ambivalence? In recent research we have shown that such situations can lead to the same information processing consequences as having discrepant automatic and deliberative attitudes.

In one study examining this notion (Petty et al., 2006, Experiment 2), us-

ing a procedure described earlier, undergraduate students were initially conditioned to like or dislike a target individual. Then, the participants received information about the target individual's attitudes on several important topics. The attitudinal information was designed to either get the person to like or dislike the target by having the target agree or disagree with the participant on several important issues. In some conditions, this information was in the same direction as the conditioning manipulation so that no attitude change would occur, and in other conditions the information was opposite in valence to the conditioning. In the latter situation, individuals rejected their earlier evaluations based on conditioning and adopted new evaluations based on the similarity information. According to the MCM, this confluence of factors produces a situation in which deliberative measures would be more sensitive to the similarity induction than would automatic measures (which will still show some evidence of the old and now rejected evaluation). However, rather than measuring automatic and deliberative attitudes following attitude change as described earlier (Petty et al., 2006, Experiment 1), in this study participants were told that the target person was a candidate for a job at their university. To evaluate the candidate, they were provided with either a strong or a weak résumé to examine. The key result was that attitudes toward the target as a job candidate were more influenced by the quality of the candidate's résumé in the condition where attitudes were changed than in conditions where attitudes were not changed. In other words, when attitudes were changed, people engaged in greater information processing as if they were attempting to resolve some ambivalence.

In other studies we have shown that when explicit attitudes change from one valence to another, people do not necessarily report feeling more ambivalent, nor do they indicate that they endorse both sides of the issue. Nevertheless, when explicit attitudes have changed in valence, people still act as if they are ambivalent by processing information on the issue to a greater extent than when attitudes are not changed. To provide more direct evidence regarding implicit ambivalence, in another study (Petty et al., 2006, Study 3) we changed the valence of participants' attitudes about a target individual or not, and then gave them an IAT to see if the target individual was more associated with doubt than when attitudes were not changed. We also measured explicit reports of conflict in this study. Consistent with the idea that an explicit change in attitudinal valence can produce implicit ambivalence (due to conflict between old rejected evaluations and newly accepted ones), participants whose attitudes were changed did not report any more explicit doubt about the target individual, but they did show more doubt on an IAT compared to when attitudes were not changed.

ATTITUDE STRENGTH

Another important feature of the MCM is the manner in which it accounts for the strength of explicit attitudes. Considerable research over the past few decades has shown that all attitudes of the same valence—when assessed with common explicit measures—are not the same in strength, where strength is defined by the extent to which the attitudes are durable (persist over time and resist change) and impactful (affect other judgments and behavior; Krosnick & Petty, 1995). There are two aspects of the MCM that relate to the strength of explicit attitudes. The first is the strength of the object–evaluation link(s) that can be approximated with measures of evaluation accessibility (see Fazio, 1995)—that is, the more accessible an object–evaluation link is, the more it will come to mind and thereby have the potential to be tapped by a deliberative attitude measure. The second is the validity tag that can be approximated with measures of evaluation confidence or certainty (see Gross et al., 1995). The more confident that people are in the validity of an object–evaluation link (depicted as a “yes” tag in Figures 1 and 2) and the more accessible this confidence is, the more they are likely to report the attitude on a deliberative measure, and the more likely this attitude is to be durable and impactful. According to the MCM, then, anything that increases the accessibility of evaluations or the confidence that one has in them will increase attitude strength.⁵

One of the most studied determinants of attitude strength emphasized by prominent attitude change theories such as the Elaboration Likelihood Model (Petty & Cacioppo, 1986), the Heuristic–Systematic Model (Chaiken, Liberman, & Eagly, 1989) and the Unimodel (Kruglanski & Thompson, 1999), is the extent of elaboration that the attitude object has received (i.e., how much thinking a person has done about it). We noted earlier that if people are doing online (downstream) thinking after an initial evaluation is accessed, it could affect the correlation between automatic and deliberative measures of attitudes because

5. Like the MCM, other theories have proposed that some factor in addition to accessibility determines the impact of one's stored attitudes. However, in these alternative frameworks, these factors typically have been defined with respect to some other concept. For example, in the accessibility–diagnosticity approach, the *diagnosticity* of an attitude is defined as the extent to which a respondent believes the attitude correctly indicates how a particular second response should be made (Feldman & Lynch, 1988). Similarly, the *relevance* of an attitude (which can impact its use) is defined only with respect to some other particular outcome variable (e.g., a particular behavior; Snyder, 1982). Thus, one speaks of the diagnosticity or relevance of an attitude for a particular outcome. In contrast, perceived validity is more general and can be assessed without reference to a specific outcome variable. Thus, it is more useful to store in memory.

new aspects of the attitude object could come to mind that conflict with (or reinforce) what is represented at the structural level (see also, Fazio & Towels-Schwen, 1999; Gawronski & Bodenhausen, 2006). However, just because thinking following retrieval of an evaluation can sometimes reduce automatic–deliberative correlations, this does not mean that evaluations that are *based on* high amounts of thinking should show lower automatic–deliberative correlations. To the contrary, when evaluations are based on high amounts of thinking, this will generally enhance the correlation between automatic and deliberative measures.

According to the MCM, attitudes based on high amounts of thinking should show higher automatic–deliberative correlations for two reasons. First, high amounts of thinking—at least if it is evaluatively congruent thinking—should produce stronger object–evaluation associations. This should enhance the impact of the automatic component on deliberative measures. Second, thinking should enhance the sense of perceived validity. This should increase the impact of the meta-cognitive component on deliberative attitudes. Indeed, the extent of elaboration has been related to attitude accessibility (Bizer & Krosnick, 2001), the extent of attitude confidence (Haugtvedt & Petty, 1992), as well as the criteria of attitude durability and impact (see Petty, Haugtvedt, & Smith, 1995, for a review). As might be expected then, the more elaboration an attitude object has received, the stronger the correlation between deliberative and automatic measures of attitudes tend to be (Nosek, 2005).

UNDERSTANDING PREJUDICE

As noted earlier, perhaps the most unique aspect of the MCM is the proposition that people can store validity tags. Besides those already mentioned, what are some of the insights that this possibility affords for understanding social psychological phenomena? Consider the domain of prejudice where many unprejudiced individuals might be expected to have both positive and negative associations to some stigmatized group, but believe that the negative link is wrong. What are the consequences of having a well-practiced “wrong” tag? Current research suggests that continually practicing a negation would not necessarily stop negative traits from coming to mind when presented with a member of the minority group (see Gawronski, Deutsch, Mbirkou, Seibt, & Strack, in press). However, according to the MCM, continual practice should lead the person to immediately think “no” or “wrong” when such stereotypes are activated. Will people who have practiced negating stereotypes be *faster* to reject them (i.e., think “wrong” after the stereotype comes to mind)

thereby correcting themselves (and perhaps others) than people who have not practiced negotiation? No research has examined this specific issue yet, but the MCM predicts that such an effect is possible.

Staying within the domain of prejudice, the MCM offers a unique perspective on various conceptualizations postulating that people are often motivated to correct for their internalized prejudice. According to several formulations (e.g., Dovidio & Gaertner, 2004; Dunton & Fazio, 1997), some White individuals have automatic negative reactions to Blacks but have egalitarian values or desires not to be prejudiced that cause them to discount their activated negativity and report positivity on explicit measures. This analysis assumes that the causal sequence is that people have preexisting attitudes that are negative, some preexisting motive to control these reactions, and these interact to determine a *constructed* deliberative positive attitude. Although this is perfectly plausible and certainly can occur, the MCM holds open another possibility—that motives can follow from preexisting positive and negative associations, with the latter negated. In other words, some people will recognize that they have both existing positive and negative associations, with the latter being unwanted. Because they find the latter to be inappropriate or wrong, they develop a motive to control these negative reactions. Thus, rather than a positive constructed attitude following from the interaction of negative automatic attitudes and a motive to control them, it could be that a motive to control negative reactions follows from the presence of both positive and negative associations to a minority group with the latter being rejected.

SUMMARY AND CONCLUSIONS

In this article we have outlined a Meta-Cognitive Model (MCM) of attitudes. In brief, the model holds that attitudes consist of evaluative associations (positive and/or negative) along with validity tags. As our discussion above implies, the MCM differs from the single, dual, and process models of attitudes in several ways. The MCM offers one integrated attitude representation (in contrast to the dual attitudes approach), but postulates that attitude objects can be linked to both positive and negative evaluations (in contrast to the single attitude approach). Viewing the attitude representation as an integrated unit rather than as separate representations activated in different situations (as advanced by some dual attitudes models) allows for joint activation of positivity and negativity in any given situation where the attitude object is encountered (assuming that people have both positive and negative associations). The possession of both positive and negative associations

can lead to explicit ambivalence when both valences are endorsed, or to implicit ambivalence when one valence is accepted and the other is rejected. We have also shown in this article how the MCM can account for a variety of attitudinal phenomena in the literature as well as make new predictions.

In terms of defining attitudes, when people have just one valence of association, the MCM reduces to the single attitude approach (unless this single valence is rejected in which case the attitude structure is more complex). Because the MCM posits an enduring attitude structure, the model lies more squarely in the camp postulating that attitudes can be stored and need not be constructed anew each time. Indeed, in the MCM, attitudes can be viewed as a person's stored evaluative associations that are not rejected. However, construction processes are necessarily involved in completing an explicit self-report if only to translate one's retrieved sense of positivity and/or negativity onto the attitude assessment offered (e.g., a bipolar scale).⁶ The MCM holds that at a minimum, when completing a deliberative scale, the attitude report is constructed from the activated evaluation(s) and any validity tags that are retrieved. At a maximum, the attitude report can also consider any other information that comes to mind and is activated by the current context. The MCM holds that explicit attitude reports will be more durable, impactful, and correlated with automatic measures (i.e., high attitude strength), when the evaluative associations are highly accessible and held with high confidence.

REFERENCES

- Barden, J., Maddux, W. W., Petty, R. E., & Brewer, M. B. (2004). Contextual moderation of racial bias: The impact of social roles on controlled and automatically activated attitudes. *Journal of Personality and Social Psychology, 87*, 5–22.
- Bassili, J. N. (1996). Meta-judgmental versus operative indices of psychological properties: The case of measures of attitude strength. *Journal of Personality and Social Psychology, 71*, 637–653.
- Betsch, T., Haberstroh, S., Molter, B., & Glöckner, A. (2004). Oops, I did it again—relapse errors in routinized decision making. *Organizational Behavior and Human Decision Processes, 93*, 62–74.
- Bizer, G. Y., & Krosnick, J. A. (2001). Exploring the structure of strength-related

6. With continued experience with bipolar scales, one integrated summary evaluation might be stored along with any positive and negative links.

- attitude features: The relation between attitude importance and attitude accessibility. *Journal of Personality and Social Psychology*, 81, 566–586.
- Bizer, G. Y., Petty, R. E., & McCaslin, M. (2005). How we conceptualize our attitudes matters: The effects of valence framing on the resistance of political attitudes. *Political Psychology*, 26, 553–568.
- Briñol, P., & Petty, R. E. (in press). Automatic and deliberative attitude change from thoughtful and non-thoughtful processes. In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures*. Mahwah, NJ: Erlbaum.
- Briñol, P., Petty, R. E. & Wheeler, S. C. (May 2003). Implicit ambivalence: Implications for attitude change. Paper presented at the Attitudinal Incongruence and Information Processing Symposium. Amsterdam, Netherlands.
- Briñol, P., Petty, R. E., & Wheeler, S. C. (2006). Discrepancies between explicit and implicit self-concepts: Consequences for information processing. *Journal of Personality and Social Psychology*, 91, 154–170.
- Byrne, D. (1961). Interpersonal attraction and attitude similarity. *Journal of Abnormal and Social Psychology*, 62, 713–715.
- Cacioppo, J.T., Gardner, W.L., & Berntson, G.G. (1997). Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review*, 1, 3–25.
- Chaiken, S., Liberman, A., & Eagly, A. H. (1989). Heuristic and systematic information processing within and beyond the persuasion context. In J. S. Uleman & J. A. Bargh (Eds.), *Unintended thought* (pp. 212–252). New York: Guilford Press.
- Cohen, J. B., & Reed, A. (2006). A multiple pathway anchoring and adjustment (MPAA) model of attitude generation and reinforcement. *Journal of Consumer Research*, 33, 1–15.
- Conrey, F.R., & Smith, E.R. (2007). Attitude representation: Attitudes as patterns in a distributed, connectionist representational system. *Social Cognition*, 25(5), 718–735.
- DeCoster, J., Banner, M. J., Smith, E. R., & Semin, G. R. (2006). On the inexplicability of the implicit: Differences in the information provided by implicit and explicit tests. *Social Cognition*, 24, 5–21.
- deLiver, Y., van der Pligt, J., & Wigboldus, D. (2007). Positive and negative associations underlying ambivalent attitudes. *Journal of Experimental Social Psychology*, 43, 319–326.
- Deutsch, R., Gawronski, B., & Strack, F. (2006). At the boundaries of automaticity: Negation as reflective operation. *Journal of Personality and Social Psychology*, 91, 385–405.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5–18.
- Dovidio, J., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). The nature of prejudice: Automatic and controlled processes. *Journal of Experimental Social Psychology*, 33, 510–540.
- Dovidio, J. F., & Gaertner, S. L. (2004). Aversive racism. In M. P. Zanna (Ed.), *Ad-*

- vances in experimental social psychology* (Vol. 36, pp. 1–51). San Diego: Academic Press.
- Dovidio, J. F., Kawakami, K., Smoak, N., & Gaertner, S. L. (in press). The roles of implicit and explicit processes in contemporary prejudice. In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures*. Hillsdale, NJ: Erlbaum.
- Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin*, *15*, 543–538.
- Eagly, A.H., & Chaiken, S. (2007). The advantages of an inclusive definition of attitude. *Social Cognition*, *25*(5), 582–602.
- Eiser, J. R., Fazio, R. H., Stafford, T., & Prescott, T. J. (2003). Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations. *Personality and Social Psychology Bulletin*, *29*, 1221–1235.
- Fazio, R. H. (1995). Attitudes as object–evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 247–283). Hillsdale, NJ: Erlbaum.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.
- Fazio, R. H. (2007). Attitudes as object–evaluation associations of varying strength. *Social Cognition*, *25*(5), 603–637.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, *54*, 297–327.
- Fazio, R. H., & Towles-Schwen, T. (1999). The MODE Model of attitude–behavior processes. In S. Chaiken & Y. Trope (Eds.), *Dual process theories in social psychology* (pp. 97–116). New York: Guilford Press.
- Feldman, J. M., & Lynch, J. G. (1988). Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of Applied Psychology*, *73*, 421–435.
- Fiske, S. T., & Pavelchak, M. A. (1986). Category-based versus piecemeal-based affective responses: Developments in schema-driven affect. In R. M. Sorrentino & E. T. Higgins (Eds.), *Handbook of motivation and cognition: Foundations of social behavior* (pp. 167–203). New York: Guilford Press.
- Gawronski, B., & Bodenhausen, G. V. (2006). Associative and prepositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*, 692–731.
- Gawronski, B., & Bodenhausen, G. V. (2007). Unraveling the processes underlying evaluation: Attitudes from the perspective of the APE Model. *Social Cognition*, *25*(5), 687–717.
- Gawronski, B., Deutsch, R. M., Mbirkou, S., Seibt, B., & Strack, F. (in press). When “just say no” is not enough: Affirmation versus negation training and the

- reduction of automatic stereotype activation. *Journal of Experimental Social Psychology*.
- Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, *46*, 107–119.
- Greenwald, A. G., & Banaji, M. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*, 4–27.
- Greenwald, A.G., McGhee, D.E., & Schwarz, J.L.K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464–1480.
- Gregg, A. P., Seibt, B., & Banaji, M. R. (2006). Easier done than undone: Asymmetry in the malleability of implicit preferences. *Journal of Personality and Social Psychology*, *90*, 1–20.
- Han, H. A., Olson, M. A., & Fazio, R. H. (2006). The influence of experimentally created extrapersonal associations on the Implicit Association Test. *Journal of Experimental Social Psychology*, *42*, 259–272
- Haugtvedt, C. P., & Petty, R. E. (1992). Personality and persuasion: Need for cognition moderates the persistence and resistance of attitude changes. *Journal of Personality and Social Psychology*, *63*, 308–319.
- Holland, R., W., Verplanken, B., & van Knippenberg, A. (2003). From repetition to conviction: Attitude accessibility as a determinant of attitude certainty. *Journal of Experimental Social Psychology*, *39*, 594–601.
- Hovland, C. I., Janis, I. L., & Kelley, H. H. (1953). *Communication and persuasion: Psychological studies of opinion change*. New Haven, CT: Yale University Press.
- Jost, J. T., Kruglanski, A. W., & Nelson, T. O. (1998). Social meta-cognition: An expansionist review. *Personality and Social Psychology Review*, *2*, 137–154.
- Kaplan, K. J. (1972). On the ambivalence–indifference problem in attitude theory and measurement: A suggested modification of the semantic differential technique. *Psychological Bulletin*, *77*, 361–372.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 774–788.
- Kehr, H. M. (2004). Implicit/explicit motive discrepancies and volitional depletion among managers. *Personality and Social Psychology Bulletin*, *30*, 315–327.
- Krosnick, J. A., Judd, C. M., & Wittenbrink, B. (2005). The measurement of attitudes. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *The handbook of attitudes* (pp. 21–76). Mahwah, NJ: Erlbaum.
- Krosnick, J. A., & Petty, R. E. (1995). Attitude strength: An overview. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 1–24). Mahwah, NJ: Erlbaum.
- Kruglanski, A. W. (1989). The psychology of being “right”: The problem of accuracy in social perception and cognition. *Psychological Bulletin*, *106*, 395–409.
- Kruglanski, A. W., & Thompson, E. P. (1999). Persuasion by a single route: A view from the unimodel. *Psychological Inquiry*, *10*, 83–109.

- Kumkale, G. T., & Albarracín, D. (2004). The sleeper effect in persuasion: A meta-analytic review. *Psychological Bulletin, 130*, 143–172.
- Lichtenstein, M., & Srull, T. K. (1985). Conceptual and methodological issues in examining the relationship between consumer memory and judgment. In L. F. Alwitt & A. A. Mitchell (Eds.), *Psychological processes in advertising effects: Theory, research and application* (pp. 113–128). Hillsdale, NJ: Erlbaum.
- Lingle, J., & Ostrom, T. M. (1981). Principles of memory and cognition in attitude formation. In R. E. Petty, T. M. Ostrom, & T. C. Brock (Eds.), *Cognitive responses in persuasion* (pp. 399–420). Hillsdale, NJ: Erlbaum.
- Maddux, W. W., Barden, J., Brewer, M. B., & Petty, R. E. (2005). Saying no to negativity: The effects of context and motivation to control prejudice on automatic evaluative responses. *Journal of Experimental Social Psychology, 41*, 19–35.
- Maio, G. R., Bell, D. W., & Esses, V. M. (1996). Ambivalence in persuasion: The processing of messages about immigrant groups. *Journal of Experimental Social Psychology, 32*, 513–536.
- Mayo, R., Schul, Y., & Burnstein, E. (2004). “I am not guilty” versus “I am innocent”: Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology, 40*, 433–449.
- Newby-Clark, I. R., McGregor, I., & Zanna, M. P. (2002). Thinking and caring about cognitive inconsistency: When and for whom does attitudinal ambivalence feel uncomfortable? *Journal of Personality and Social Psychology, 82*(2), 157–166.
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General, 134*, 565–584.
- Osgood, C. E., Suci, G. J., & Tanenbaum, P. H. (1957). *The measurement of meaning*. Urbana, IL: University of Illinois Press.
- Petty, R. E. (2006). A metacognitive model of attitudes. *Journal of Consumer Research, 33*, 22–24.
- Petty, R. E., & Briñol, P. (2006). A meta-cognitive approach to “implicit” and “explicit” evaluations: Comment on Gawronski and Bodenhausen (2006). *Psychological Bulletin, 132*, 740–744.
- Petty, R. E., & Briñol, P. (in press). Implicit ambivalence: A meta-cognitive approach. In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures*. Mahwah, NJ: Erlbaum.
- Petty, R. E., Briñol, P., & Tormala, Z. L. (2002). Thought confidence as a determinant of persuasion: The self-validation hypothesis. *Journal of Personality and Social Psychology, 82*, 722–741.
- Petty, R. E., Briñol, P., Tormala, Z. L., & Wegener, D. T. (2007). The role of meta-cognition in social judgment. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 254–284). New York: Guilford Press.
- Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. New York: Springer/Verlag.
- Petty, R. E., Haugtvedt, C. P., & Smith, S. M. (1995). Elaboration as a determinant

- of attitude strength: Creating attitudes that are persistent, resistant, and predictive of behavior. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 93–130). Mahwah, NJ: Erlbaum.
- Petty, R. E., & Jarvis, W. B. G. (1998, October). *What happens to the old attitude when attitudes change?* Presented at the annual meeting of the Society for Experimental Social Psychology, Lexington, KY.
- Petty, R. E., Tormala, Z. L., Briñol, P., & Jarvis, W. B. G. (2006). Implicit ambivalence from attitude change: An exploration of the PAST Model. *Journal of Personality and Social Psychology, 90*, 21–41.
- Petty, R. E., Tormala, Z. L., & Rucker, D. D. (2004). Resisting persuasion by counterarguing: An attitude strength perspective. In J. T. Jost, M. R. Banaji, & D. A. Prentice (Eds.), *Perspectivism in social psychology: The yin and yang of scientific progress* (pp. 37–51). Washington, DC: American Psychological Association.
- Petty, R. E., Wells, G. L., & Brock, T. C. (1976). Distraction can enhance or reduce yielding to propaganda: Thought disruption versus effort justification. *Journal of Personality and Social Psychology, 34*, 874–884.
- Petty, R. E., Wheeler, S. C., & Tormala, Z. L. (2003). Persuasion and attitude change. In T. Millon & M. J. Lerner (Eds.), *Handbook of psychology: Vol. 5: Personality and social psychology* (pp. 353–382). Hoboken, NJ: John Wiley and Sons.
- Plant, E. A., & Devine, P. (1998). Internal and external motivation to respond without prejudice. *Journal of Personality and Social Psychology, 75*, 811–832.
- Priester, J. R., & Petty, R. E. (2001). Extending the bases of subjective attitudinal ambivalence: Interpersonal and intrapersonal antecedents of evaluative tension. *Journal of Personality and Social Psychology, 80*, 19–34.
- Rosenberg, M. (1965). *Society and the adolescent self-image*. Princeton, NJ: Princeton University Press.
- Rydell, R. J., & McConnell, A. R. (2006). Understanding implicit and explicit attitude change: A systems of reasoning analysis. *Journal of Personality and Social Psychology, 91*, 995–1008.
- Rydell, R. J., McConnell, A. R., Mackie, D. M., & Strain, L. M. (2006). Of two minds: Forming and changing valence-inconsistent implicit and explicit attitudes. *Psychological Science, 17*, 954–958.
- Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition, 25*(5), 638–656.
- Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Blackwell handbook of social psychology: Intraindividual processes* (pp. 436–457). Malden, MA: Blackwell Publishers.
- Sherman, J. W. (in press). Controlled influences on implicit measures: Confronting the myth of process-purity and taming the cognitive monster. In R. E. Petty, R. H. Fazio, & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures*. Hillsdale, NJ: Erlbaum.
- Snyder, M. (1982). When believing means doing: Creating links between attitudes and behavior. In M. P. Zanna, E. T. Higgins, & C. P. Hermann (Eds.),

- Consistency in social behavior: The Ontario Symposium* (Vol. 2, pp. 105–130). Hillsdale, NJ: Erlbaum.
- Thurstone, L. L. (1928). Attitudes can be measured. *American Journal of Sociology*, 33, 529–554.
- Tormala, Z. L., DeSensi, V. L., & Petty, R. E. (2007). Resisting persuasion by illegitimate means: A meta-cognitive perspective on minority influence. *Personality and Social Psychology Bulletin*, 33, 354–367.
- Van Overwalle, F., & Siebler, F. (2005). A connectionist model of attitude formation and change. *Personality and Social Psychology Review*, 9, 231–274.
- Vargas, P. T. (2004). The relationship between implicit attitudes and behavior: Some lessons from the past and directions for the future. In G. Haddock & G. R. Maio (Eds.), *Contemporary perspectives on the psychology of attitudes* (pp. 275–297). New York: Psychology Press.
- Wells, G. L., Olson, E. A., & Charman, S. D. (2003). Distorted retrospective eyewitness reports as functions of feedback and delay. *Journal of Experimental Psychology: Applied*, 9, 42–52.
- Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. L. Martin & A. Tesser (Eds.), *The construction of social judgments* (pp. 37–65). Hillsdale, NJ: Erlbaum.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101–126.
- Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, 81, 815–827.
- Zanna, M. P., & Rempel, J. K. (1988). Attitudes: A new look at an old concept. In D. Bar-Tal & A. W. Kruglanski (Eds.), *The social psychology of knowledge* (pp. 315–334). Cambridge, UK: Cambridge University Press.