

## THE ITERATIVE REPROCESSING MODEL: A MULTILEVEL FRAMEWORK FOR ATTITUDES AND EVALUATION

William A. Cunningham  
*Ohio State University*

Philip David Zelazo  
*University of Minnesota*

Dominic J. Packer  
*Ohio State University*

Jay J. Van Bavel  
*Ohio State University and University of Toronto*

Dual-process models of attitudes highlight the fact that evaluative processes are complex and multifaceted. Nevertheless, many of these models typically neglect important interactions among processes that can contribute to an evaluation. In this article, we propose a multilevel model informed by neuroscience in which current *evaluations* are constructed from relatively stable *attitude* representations through the iterative reprocessing of information. Whereas initial iterations provide relatively quick and dirty evaluations, additional iterations accompanied by reflective processes yield more nuanced evaluations and allow for phenomena such as ambivalence. Importantly, this model predicts that the processes underlying relatively automatic evaluations continue to be engaged across multiple iterations, and that they influence *and* are influenced by more reflective processes. We describe the Iterative Reprocessing Model at the computational, algorithmic, and implementational levels of analysis (Marr, 1982) to more fully characterize its premises and predictions.

Recent advances in neuroscientific methods have provided researchers with an unprecedented opportunity to examine the neural correlates of

---

Please address correspondence to William A. Cunningham, Department of Psychology, The Ohio State University, 100G Lazenby Hall, 1827 Neil Avenue, Columbus, OH 43210. E-mail: [cunningham.417@osu.edu](mailto:cunningham.417@osu.edu).

human thought and emotion. Inspired by achievements in cognitive domains, social psychologists have recently turned to neuroscience to add a layer of understanding to the processes of social judgment and behavior. Efforts to understand social processes across multiple levels of analysis have inspired a new field, variably called social neuroscience (Cacioppo, Berntson, Sheridan, & McClintock, 2000), social cognitive neuroscience (Ochsner & Lieberman, 2001), and the social brain sciences (Adolphs, 2003). Incorporating knowledge regarding brain function into our understanding of attitudes and evaluation promises to lead to the refinement of theoretical models and generation of novel hypotheses. We propose an initial multilevel framework for understanding some of the core operating characteristics of the human evaluative system. We contend that consideration of the neural and computational processes underlying evaluation provides several novel and provocative answers to outstanding “controversies” among attitude theorists, including whether or not attitudes are constructed or stable, whether or not there is one ‘true’ attitude, the relation between “automatic” and “controlled” evaluations, and the nature of attitudinal ambivalence. We also believe that our model of the human evaluative system will inform other debates in the field of attitudes and related fields, such as prejudice, judgment and decisionmaking, and emotions.

Our discussion will be grounded in the Iterative Reprocessing (IR) Model of evaluation (Cunningham & Zelazo, 2007). Whereas most animals’ evaluations rely primarily on automatically activated, nonconscious associations between a stimulus and an outcome/response, human beings’ evaluations are complemented by additional reflective processes, supported by prefrontal cortical (PFC) brain regions. The IR Model highlights the potential interactions between relatively automatic and reflective processes (Cunningham & Zelazo, 2007; Packer, Kesek, & Cunningham, in press). The fundamental assumption underlying our model is that brain systems are organized hierarchically, such that lower-order automatic processes influence *and* are influenced by higher-order, reflective processes (see Zelazo & Cunningham, 2007). Importantly, reflective processes do not supplant automatic ones; rather, these processes work in concert in order to generate dynamic evaluations that incorporate prior associational histories of stimuli with additional information, including the context and current goals. Increasingly reflective processes allow for the more nuanced and/or goal-congruent evaluations of stimuli needed to navigate complex environments, self-regulate, and appraise abstract concepts (e.g., liberty, equality, fraternity).

In the following sections, we describe the IR Model at each of Marr’s (1982) three levels of analysis for understanding complex information

processing systems. First, the *computational* level concerns what evaluative processing accomplishes. Second, the *implementational* level examines how the brain processes information (the neural manifestation of computations). Third, the *algorithmic* level describes how information is represented and processed. Parallels among all three levels of analysis are highlighted. Taken together, these levels of analysis provide a more comprehensive understanding of some of the core operating characteristics of the human evaluative system. Following this discussion, we apply the IR Model to several outstanding controversies in the study of attitudes.

### **THE ITERATIVE REPROCESSING MODEL: THE COMPUTATIONAL LEVEL**

According to the IR Model, an individual's relatively stable stored representations regarding the valence of a stimulus are conceptually distinct from current evaluation of the stimulus (Cunningham & Zelazo, 2007; Cunningham, Packer, Kesek, & Van Bavel, in press; see also Eagly & Chaiken, this issue). *Evaluative processes* generate evaluations by integrating activated preexisting attitudes with additional information about the stimulus, context, and goal states. Thus, whereas an *attitude* is a relatively stable set of representations (only some of which may be active at any time) stored in memory, an evaluation reflects the current processing state of the entire evaluative system (which is influenced by those aspects of the attitude that are currently active).

To understand the mechanisms involved in the processing of attitudes and the generation of evaluations, it is important to note that evaluations are the dynamic result of an integrated set of distributed processes, each of which responds to and resolves specific computational problems (see Cunningham & Johnson, 2007). In other words, rather than corresponding to a discrete neural process, evaluation is an emergent property of multiple processes that unfold over time. Although the IR Model assumes that the networks involved in evaluation are highly distributed and dynamic, because these networks are hierarchical, there is a common set of processes that are consistently involved in generating current evaluations. Specifically, *lower-order* evaluative processes continue to provide information about the valence (positivity and negativity) and the arousal value (high vs. low) of stimuli, even as *higher-order* reflective processes come online.

Encountered or imagined stimuli (e.g., people, objects, or abstract concepts) initiate an iterative sequence of processes that interpret and then reinterpret these stimuli in light of an increasingly wide range of consid-

erations (see Figure 1). The IR Model assumes that evaluative processes are part of an *iterative cycle*: with every iteration, the current evaluation of a stimulus can be adjusted in light of additional contextual and motivational information in order to create an updated evaluation in line with finer stimulus detail, the context, and/or current goals. Information is continually passed back from relatively higher-order to relatively lower-order processes and the evaluation is recalculated. This “reseeding” of information allows for the foregrounding of relevant (and backgrounding of irrelevant) attitude representations and contextual information in order to develop a more nuanced evaluation congruent with current goals.

In general, each iteration of the cycle allows for additional reflective processing<sup>1</sup>. Thus, the IR Model implies a rough continuum from “automatic” evaluations, entailing few iterations and a limited set of cognitive operations, to more “reflective” evaluations, entailing more iterations and cognitive operations (see Cunningham & Johnson, 2007). Initial iterations generally involve processing in subcortical, limbic regions and give rise to rapid evaluations based on relatively simple preexisting associations or innate preferences (LeDoux, 1996; Phelps, 2006; Whalen, 1998). With the passage of time, more reflective processes (mediated by the PFC) come online, providing an opportunity for the reinterpretation and modulation of earlier evaluations (e.g., Cunningham, Johnson, Gateby, Gore, & Banaji, 2003; Cunningham, Raye, & Johnson, 2004). Consistent with a hierarchical view of brain systems (e.g., Bunge & Zelazo, 2006; Crone & von der Molen, 2004; Luria, 1966), the IR Model proposes that these higher-order processes continually reseed information to lower-order processes in order to generate updated evaluations. At each iteration, the current evaluation serves as input for ongoing evaluative processing; as such, earlier evaluations are likely to bias subsequent evaluative processing by influencing attention, information seeking, stimulus construal, etc. Importantly, while conscious deliberation exerts an influence on evaluative processing, information about the valence and arousal value of a stimulus continues to be represented in subcortical brain structures.

An example can help to illustrate the operation of the IR Model. Imagine that someone is standing next to a sullen looking young man of Arab descent on the subway. Primed by television coverage of terrorist threats, the demeanor of the man coupled with his ethnicity may auto-

---

1. It is important to note that additional time on task does not necessitate more reflective processing. People can ruminate about a stimulus for a long time before thinking about it in a different way.

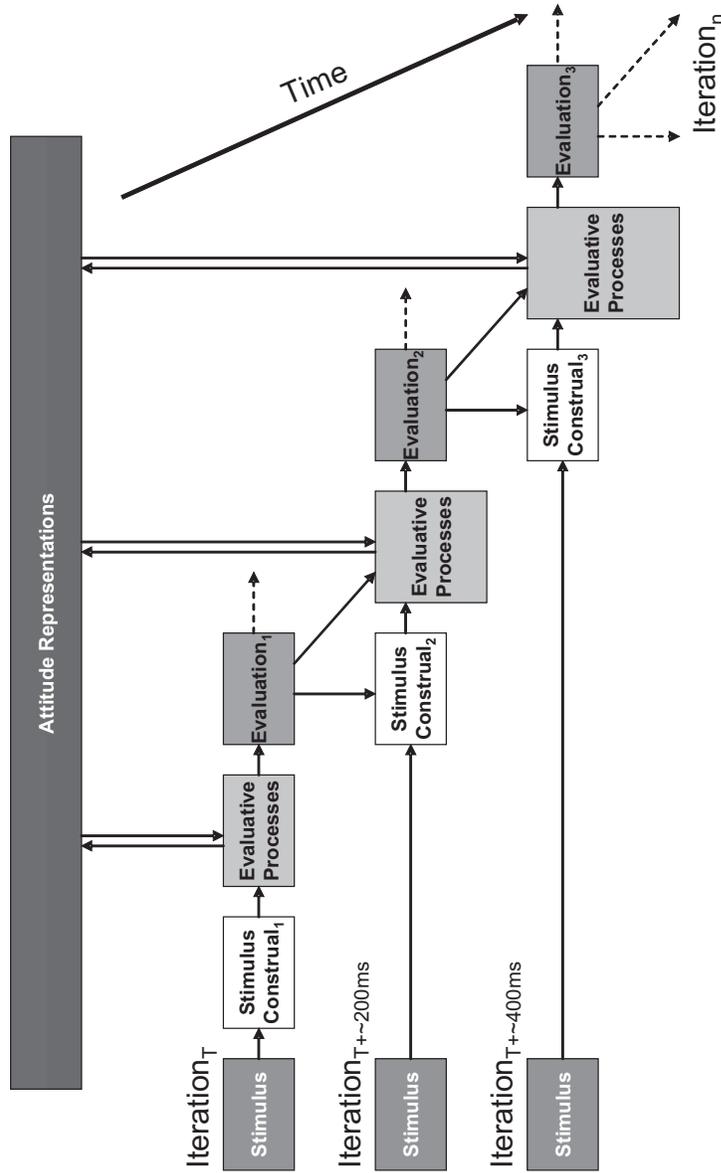


FIGURE 1. The evaluative cycle. At each iteration, evaluative processes rely on a particular construal of the stimulus and retrieve attitude representations to generate an evaluation. This evaluation can be used to influence the next iteration of evaluative processing, direct behavior, or both. With each additional iteration, the complexity of evaluative processing can increase.

matically activate stored representations associated with terrorists, giving rise to a rapid negative evaluation. This initial evaluation may bias subsequent iterations of evaluative processing, including interpretations of the man's behavior; if the young man reaches for his wallet, the subway rider may worry that he is about to pull out a gun, and may indeed mistakenly perceive the wallet as a weapon (Payne, 2001). At some point, however, the perceiver will likely detect a mismatch between his or her current evaluation of the target and reality—for example, the perceiver may see that what her or she thought was a gun is actually a wallet and that the man is about to give money to a homeless person. This discrepancy will trigger additional iterations and increase reflective processing to resolve the conflict between the current evaluation and reality. These iterations may involve a search for and retrieval of positively valenced representations of young Arab men or an effort to suppress previously activated negative representations; strategies that involve reprocessing and reinterpreting the man and the perceiver's reactions to him.

As this example illustrates, reflective processes allow for more nuanced evaluations, integrating activated attitude representations with additional information, including information regarding the social context. According to the IR Model, evaluation is not the result of a single or dual process occurring within a fixed time frame. Some judgments may be reached rapidly and remain stable across the lifespan, whereas others may be continually altered and updated every few hundred milliseconds as new information is processed. The extent to which the evaluation of any particular stimulus is reprocessed (i.e., the number of iterations it receives) is likely to depend on a host of personal and situational factors, including differences in cognitive ability, motivation, and opportunity (Fazio, 1990), as well as PFC functionality (e.g., Bechara, 2004; Zelazo, 2004).

Reflective processes take time and effort, and it is inefficient, if not impossible, to process the myriad intricacies of all stimuli and events—some simplification is always needed, and is needed more for some evaluations than for others. To understand when people are more likely to generate complex evaluations, the IR Model proposes that two competing motivational drives influence the extent of evaluative processing across situations. First, a drive to minimize the discrepancy between one's evaluation and the hedonic environment (i.e., to minimize error) increases reflective processing during evaluation. Second, a drive to minimize processing demands decreases in reflective processing during evaluation. These opposing drives create a dynamic tension that can help individuals to strike a delicate balance between an initial "gut" response and evaluations that are more nuanced, but not computationally

catastrophic. The influence of these competing motivations likely varies as a function of situational demands, current goals, and individual differences in processing style. These motivations represent two of the core goals of evaluative processing, highlighting the computational aspect of the IR Model (i.e., what evaluative processes accomplish).

### NEURAL CORRELATES: THE IMPLEMENTATIONAL LEVEL

Building on recent advances in cognitive and affective neuroscience, we present a preliminary neural model of the networks involved in evaluation (see Cunningham & Zelazo, 2007; Zelazo & Cunningham, 2007 for more details). The IR Model of evaluation is predicated on the interaction between different neural component processes. According to the IR Model, evaluations are the result of a series of iterations during which information about a stimulus is reprocessed and integrated with new information about the stimulus and the context in which it occurs.

The proposed neural mechanisms underlying iterative reprocessing are shown in Figure 2. At the heart of the model are brain systems involved in the generation and processing of affect. Specifically, subcortical brain regions such as the amygdala and the ventral striatum (especially the nucleus accumbens) are well situated for the rapid automatic evaluation of stimuli—based on innate biases (e.g., Öhman & Mineka, 2001) and learning (e.g., Armony & Dolan, 2002; Davis, 1997; Phelps et al., 2001; Whalen, 1998). For example, a direct route to the amygdala through the thalamus (i.e., the “low road” to amygdala activation), may permit evaluative processing before a stimulus is even fully processed by cortical sensory areas (LeDoux, 1996). Thus, when a stimulus is encountered, information about the stimulus is fed (via the direct, subcortical route) from the thalamus to the amygdala triggering an unreflective motivational tendency to approach or avoid the stimulus, and producing a series of physiological responses and reflexive reactions that are mediated by the hypothalamus (among other regions). These relatively undifferentiated physiological responses prepare the body for immediate action—fighting or fleeing—while additional neural processes continue to disambiguate the motivational and evaluative implications of the stimulus. Importantly, although subcortical structures support automatic evaluations, we suggest that these regions remain involved throughout the evaluative cycle. In other words, even as additional reflective processes are incorporated in evaluation processing, subcortical processes may maintain an ongoing and important role in generating current evaluations.

Following this initial evaluation, the physiological response is event-

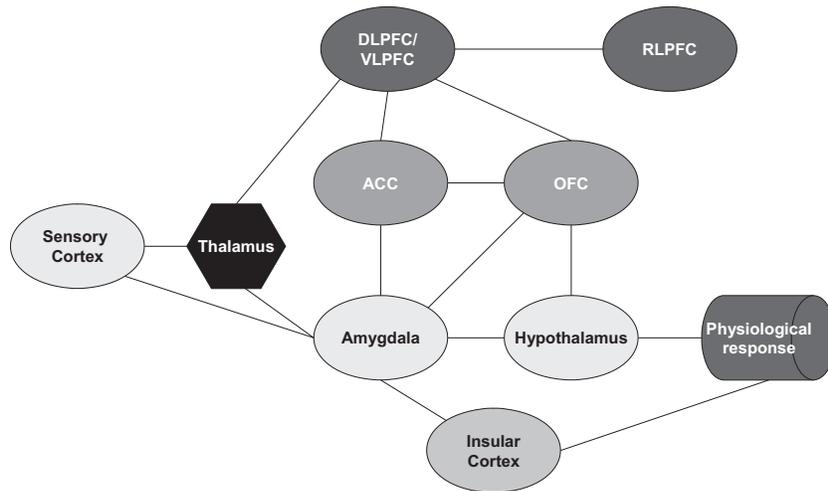


FIGURE 2. Neural circuitry underlying evaluation. A simplified neural model of the Iterative Reprocessing model is presented. Links between regions discussed in the article are denoted by solid lines. We acknowledge that not all anatomical links are represented. Information about a stimulus is processed by the thalamus and projected to the amygdala, leading to an initial evaluation that is associated with a tendency to approach or avoid the stimulus. Additional iterations can also include processing by the insula, orbitofrontal cortex (OFC), and anterior cingulate cortex (ACC)—as well as more detailed sensory processing. Visceral changes following evaluation are guided by the hypothalamus and other regions associated with autonomic control. Additional recruitment of the prefrontal cortex, especially regions of the ventrolateral prefrontal cortex (VLPFC), dorsolateral prefrontal cortex (DLPFC), and rostralateral prefrontal cortex (RLPFC), may serve a role in strategic reprocessing of stimuli and may serve to regulate evaluative processing by amplifying or suppressing attention to certain aspects of the stimulus or situation.

ally registered in the somatosensory cortex (and the insular cortex in particular) allowing for the representation of information about the current bodily state. Given direct connections between the insula and amygdala, the represented bodily state can be integrated into subsequent iterations of evaluative processing (Damasio, 1994; 1996, Critchley, Wiens, Rotshtein, Öhman, & Dolan, 2004; see also Rolls, 2000). During these subsequent iterations, information from the thalamus, more detailed information from the sensory cortices, and the bodily state all provide input into the next evaluation. These multiple inputs to the amygdala can serve to generate a more complex evaluation within a

time period still typically labeled “automatic”—that is, *several* initial iterations may all occur within a few hundred milliseconds of stimulus perception (Oya, Kawasaki, Howard, & Adolphs, 2002).

As information about the stimulus is reprocessed, more reflective processes may influence the evaluation.<sup>2</sup> After the initial iterations, amygdala projections to the orbitofrontal cortex (OFC) allow for a comparison of expected rewards and punishments with current experience. Whereas the subcortical systems provide a low-resolution estimate of likely outcomes, regions of the OFC may be involved in integrating amygdala output with current experience, allowing the current context to play a role in shaping the evaluation (e.g., Blair, 2004; Beer, Heery, Keltner, Scabini, & Knight, 2003; Rolls, 2000; Rolls, Hornak, Wade, & McGrath, 1994). The OFC receives input from multiple sensory modalities and may play an important role in providing a common metric for representing and comparing disparate aspects of evaluative information (Rolls, 2000). Further, the OFC provides a degree of top-down regulation of amygdala activation, and it can also affect physiological responses to a stimulus via projections to the hypothalamus.

In many cases, the affective responses to the situation mediated by the amygdala and OFC will yield a motivationally sufficient evaluation. In other cases, however, this joint processing will yield too much residual uncertainty or evidence of conflict, as when the stimulus is ambivalent, not providing expected rewards, etc. According to the IR Model, this conflict triggers anterior cingulate cortex (ACC) activation, which may then initiate additional reprocessing of the stimulus in the lateral PFC (Bunge & Zelazo, 2006; Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004). ACC activation (especially the dorsal ACC) is likely involved in conflict monitoring and detection (e.g., Bush, Luu, & Posner, 2000; Carter et al., 1998), whereas reprocessing by lateral regions of the PFC likely corresponds to increasingly conscious control (e.g., MacDonald, Cohen, Stenger, & Carter, 2000). Additional reflection proceeds, as needed, through increasingly higher-order regions within the hierarchy of the PFC: from ventrolateral (VLPFC) to dorsolateral (DLPFC) to rostralateral (RLPFC) areas. More lateral PFC-mediated processing allows evaluation to be regulated in a top-down fashion by deliberately amplifying or suppressing attention to certain aspects of the stimulus, and hence changing processing on subsequent iterations. This iterative reprocessing will not necessarily generate altogether new evaluations (let alone new attitudes), but it will likely modulate the current

---

2. It is important to note that more iterations do not necessarily mean more reflection; rather, they just provide the opportunity for more reflection.

evaluation by modulating activity in lower-order regions (e.g., Cunningham, Johnson et al., 2004; Ochsner, Bunge, Gross, & Gabrieli, 2002; Ochsner et al., 2004).

More complex networks of evaluative processing permit more complex construals of a stimulus, in part because additional information can be integrated into the construal during each iteration, and in part because these networks support the formulation and use of higher-order rules for selecting certain aspects of a stimulus to which to attend (Bunge & Zelazo, 2006). Top-down PFC processes foreground specific aspects of information (and background others), and these reweighted representations are used to reseed initial evaluative processing by influencing ongoing perception and processing of the stimulus. The PFC may play an important role in foregrounding not only activation patterns, but also in making evaluative processing consistent with system-level motivations to minimize error and/or minimize processing complexity, as a function of current goals and contextual demands/constraints (e.g., Cunningham, Raye, & Johnson, 2005; Cunningham, Van Bavel, & Johnsen, in press). This characterization of the PFC is consistent with its hypothesized role in allowing for higher levels of reflective consciousness via reprocessing (Zelazo, 2004) and in the monitoring and control of cognition and behavior (e.g., Carver & Scheier, 2001; Shallice, 1982; Stuss & Benson, 1986). Taken together, the dynamic interactions between different brain regions support the flexible and complex process of evaluation.

### **THE NATURE OF REPRESENTATIONS: THE ALGORITHMIC LEVEL**

The distinction between *attitudes* as relatively stable representations and evaluations as the current state of evaluative processing raises important questions about how attitudes are transformed into evaluations. To address this issue, we take a connectionist approach to representation in which both the attitudes and evaluations can be thought of as distributed patterns (see also Conrey & Smith, this issue). Following connectionist models of memory, attitudes are represented as the preexisting connection weights between units, whereas evaluations reflect the current activation of the units themselves (O'Reilly & Munakata, 2000). Thus, the relation between attitudes and evaluations is relatively straightforward—stronger connection weights increase the probability that similar patterns of activity will follow from the perception of an internally or externally generated stimulus. Importantly, changes in stimulus construal or top-down foregrounding can trigger activation

patterns different from those that are automatically activated during initial iterative cycles.

This conceptualization allows for both stability and flexibility within the evaluative system—two properties that are essential for adaptive function across situations. Specifically, connection weights update relatively slowly to represent statistically reliable knowledge extracted from the environment (i.e., attitudes). Although evaluative activation patterns are partially influenced by connection weights, these patterns can shift rapidly as a function of context, goals, or the activation of additional representations. As the activation pattern shifts, the current evaluation is modified. This model requires only a single set of representations to account parsimoniously for the properties of *both* “fast” and “slow” learning systems, typically reported in social cognition studies (see Chaiken & Trope, 1999).<sup>3</sup>

Within this model, the evaluative system extracts meaningful information from the weights to generate an evaluation vis-à-vis an activation pattern. We further propose that evaluative information is extracted from activation patterns through processes similar to factor analysis.<sup>4</sup> Factor analysis extracts variance from a dataset by distinguishing the most reliable patterns of data from the less reliable patterns. In this sense, the goal of evaluation is to extract the predictive eigenvalues or latent variables from a much larger set of information. Because these latent variables are, by definition, generated models of the full set of information, they have some degree of discrepancy from the original input. Thus, evaluations are rarely a perfect replication of the underlying attitudes (which are, in turn, imperfect replications of the real world).

In their classic book on the measurement of meaning, Osgood, Suci, and Tannenbaum (1967) found that the most dominant factor within any pattern of semantic information was valence. Following this insight, we predict that an estimate of valence (positivity to negativity) may be rapidly extracted from patterns of neural activity triggered by a stimulus. Forcing a single-factor solution on an activated pattern has advantages

---

3. Although we suggest one set of representations for the model, we should be clear that this is not to suggest that memories are all stored in one location or through one process. Rather, we suggest that the representations that allow for evaluations are potentially accessible throughout the evaluative system.

4. Given the conceptual overlap between connectionist frameworks and factor analysis, for the remainder of this section we will translate the conceptual points into the more familiar factor analysis framework in order to convey them more clearly. We acknowledge that the mapping is imperfect, but we suggest that they are similar enough to preserve the important information. Factor analysis provides a useful metaphor for understanding the algorithmic level of the evaluative system.

for evaluative processing—namely, it is quick, efficient, simple, and may provide a useful heuristic for guiding behavior. In terms of processing tradeoffs, a single-factor solution reduces computational demands at the cost of increasing prediction errors. A stimulus is likely to activate a stored pattern from the connection weights (an attitude), and this simple mathematical computation can easily convert this activated attitude into an evaluation. Thus, in our framework, this rapidly computed evaluation is similar to the concept of an “implicit” attitude (or summary attitude), and can occur during the first iteration(s) of evaluative processing.

Although this first iteration of information processing is similar to the automatic processes described by traditional dual attitude models, the IR Model differs from these models in terms of subsequent processing. Traditional dual attitude models (e.g., Devine, 1989; Greenwald & Banaji, 1995; Rydell & McConnell, 2006; Wilson, Lindsey, & Schooler, 2000) suggest that following automatic attitude activation, behavior is determined by this automatic attitude unless controlled processing inhibits the automatic attitude and replaces it with a different explicit attitude representation. In contrast to this view, the IR Model proposes that the very same attitude weights and pattern of activation generated in the first iterative cycle are used in subsequent evaluative processing. Specifically, within a few hundred milliseconds of initial processing, another iteration of the processing cycle begins updating the evaluation. At this stage, additional processes can influence the evaluation. For example, the physiological response of the body from a preceding iteration can serve as one input into the next iteration. Several iterations may occur before one is consciously *aware* of their current evaluation, and perhaps an additional iteration is necessary before that person is then able to apply volitional control to their evaluation. Indeed, according to the model, consciousness is a multistep process in which awareness emerges in a graded fashion over the first few iterations; a particular level of awareness is required before one can compare one’s subjective impression of his or her current evaluation with his or her goals (see Zelazo & Cunningham, 2007 for more details regarding levels of consciousness). If motivated to do so, individuals may then engage in top-down, controlled processing to shape their evaluations in subsequent iterations. Importantly, these controlled processes then interact with lower-order systems to reseed or shape the current activation patterns.

Independent of top-down control, one of the most important things that can occur as processing continues is the extraction of additional patterns or factors from the total set of information. For example, information regarding some stimuli may contain patterns that can be

simultaneously represented as positive *and* negative (i.e., an ambivalent attitude; see Cacioppo & Berntson, 1994; Priester & Petty, 1996). In the first iteration, this ambivalent information would likely be collapsed as a neutral point on a single bipolar valence factor. However, lower-order prefrontal processes (perhaps implemented in the OFC) recruited in subsequent iterations likely allow for multifactor patterns to be computed and processed simultaneously. We also speculate that factors may eventually be represented hierarchically, with latent factors-of-factors represented at the highest level of abstraction.

In terms of overarching processing goals, extracting more factors reduces the error (i.e., increases the amount of information explained by the current evaluation), but also increases computational demands. This additional processing can, however, compromise the ability of the evaluative system to come to an unambiguously positive or negative response. For example, one could extract a factor representing the dishonesty of a political candidate and another representing the candidate's position on a heated political issue. If dishonesty and the position are both evaluated negatively, the two factors do not conflict and processing can continue. Indeed, a one-factor evaluation would likely follow—negativity. On the other hand, if the candidate's position is evaluated positively (e.g., because it is consistent with one's own ideology) conflict will exist in the system—both positive and negative information are currently active. To the extent that conflict is generated and detected *and* an individual requires a less ambiguous evaluation or a binary behavioral choice (i.e., to approach vs. avoid), the conflict must be resolved. Although conflict resolution is clearly necessary in the case of discrepantly valenced information, it may also be required when different activated features are of the same valence, as when selecting among attractive alternatives (Festinger, 1957). Thus, processes involved in the weighting of information, selecting among features, and foregrounding of relevant information are needed. In terms of the factor analysis metaphor, second-level factors may subsume first-level factors in a hierarchal fashion. These second-level factors can serve to enhance some activation patterns over others and allow for more abstract resolution of conflict. For example, when a Democrat is conflicted about his marriage to a Republican (or *visa versa*), he can re-represent the relationship at a higher-level of analysis—they both share a passion for the American political system. As the computations become increasingly complex, additional explanatory factors can be created to organize and make sense of the factors at the first level. We suggest that the hierarchical structure of PFC regions may be particularly relevant for this task. Importantly, these different processes are operating on the same broad set of information (weights), and are simply using factors to extract,

organize, and re-represent this same information in different ways and at different levels of analysis.

## DISCUSSION

We propose that theories of attitudes and evaluations can be informed by taking a multilevel approach that emphasizes the processes, neural architecture, and computations of the evaluative system. In the following sections, we address how aspects of the IR Model can inform our understanding of several outstanding controversies in attitude theory.

## CONSTRUCTION AND STABILITY

A longstanding debate in the attitude literature concerns the extent to which attitudes are stored in memory versus constructed as a function of context and current goals. A foundational assumption for most researchers in the field is that attitudes are based on relatively stable memory representations that, to a greater or lesser degree, guide behavior across space and time (Fazio, this issue; Petty, Briñol, & DeMarree, this issue). In contrast, constructivist accounts highlight the sensitivity of attitudes to context and elaboration (Schwarz, this issue; Wilson & Hodges, 1992). If attitudes are stable stored representations applied to perceived stimuli, how can they change from situation to situation?

Interestingly, the IR Model predicts and necessitates both stability and flexibility, and as such, both representation and construction. If *attitudes* are relatively stable memory *weights* in a connectionist evaluative system and *evaluations* represent the current processing state of the evaluative system, or the current *activation pattern* in the evaluative system, there is no problem with positing both hypotheses (see also Eagly & Chaiken, this issue for a similar taxonomy). In other words, the dynamic nature of evaluations and their responsiveness to contextual and motivational information is not a challenge to the concept of stable representation in that these evaluations are constructed from preexisting attitudes (in the form of memory traces).

## HOW MANY ATTITUDES?

By distinguishing the concepts of *attitudes* from *evaluations*, the IR Model suggests that any given set of weights (attitude) can give rise to multiple patterns of activation (evaluations). In this sense, a single attitude (defined as the full set of stored information about the valence of a stimulus)

can give rise to a multitude of different evaluations. Evaluative processes mediate the relation between attitudes and evaluation, and play an important role in the number and type of evaluations that emerge. The factor analytic metaphor presented earlier suggests that a first pass processing on a set of weights may lead a single factor analogous to a summary attitude. However, it is important to note that the summary attitude or first factor does not necessarily capture the full extent of information that one has available about the stimulus. Rather, varying strengths of connection weights may increase the accessibility of a particular type of evaluation (i.e., positive or negative) causing it to be activated relatively automatically across time and situations (thus appearing fairly stable). However, reflective processing can permit nearly infinite activation patterns, which may draw upon additional stored attitudinal information, as well as integrate additional information regarding context and goals.

#### NATURE OF “IMPLICIT” AND “EXPLICIT” ATTITUDES

Many current models of attitudes are based on a dual-attitude framework. These models typically highlight the distinction between evaluations that operate relatively automatically (and perhaps unconsciously) and those that are the result of more deliberate conscious processing (see Rydell & McConnell, 2006). Following this distinction, over the past two decades, research has consistently found that attitudes measured automatically may be of a different valence than attitudes measured with self-report scales (Nosek, 2005). For example, people show dramatically more preference for in-groups (e.g., White, American, Christian, etc. in our typical samples) over out-groups on “implicit” measures (e.g., Bonafide Pipeline—Fazio, Jackson, Dunton, & Williams, 1995; Implicit Association Test—Greenwald, McGhee, & Schwartz, 1998; Affect Misattribution Paradigm—Payne, Cheng, Govorun, & Stewart, 2005) than on self-report scales, where university samples typically report strong egalitarian attitudes (Cunningham, Nezlek, & Banaji, 2004; Nosek, Banaji, & Greenwald, 2002). These data have been taken by some to suggest that attitude representations may be dissociated—implicit measures tap an implicit representation whereas explicit measures tap separate explicit representations.

Although these dual-attitude models have highlighted the fact that evaluative processes are potentially complex and multifaceted, we suggest that the notion of a strict dissociation between implicit and explicit processes/representations is unlikely to be accurate (see also Strack & Deutsch, 2004). In contrast, the IR Model proposes that evaluations arise

from multiple cognitive and affective processes that work in concert, rather than separately and independently, to make evaluations about the world. Thus, for the most part, any given process does not typically work alone in an all-or-none fashion, but rather various combinations of processes generate qualitatively different evaluations. With this in mind, a single set of attitude representations can produce quite different evaluations depending on the number of iterations and the nature of the processes brought to bear on the judgment.

The IR Model thus proposes that there is a continuum from relatively automatic to controlled evaluative processes that can operate on a set of attitude representations. At the early stages of processing, the strongest weights associated with an attitude give rise to a specific pattern of activation, and result in quick and automatic evaluations.<sup>5</sup> However, with more iterations and the potential for reflective processing in areas of the PFC, evaluations are shaped by a dynamic interaction of several bottom-up and top-down processes. This interaction allows for the foregrounding and backgrounding of particular patterns of activation in accordance with current contexts and goals. Whereas evaluations that are based on a small number of iterations of the evaluative cycle are relatively automatic, in that they are obligatory and may occur without conscious monitoring, evaluations based on additional iterations are generally more reflective. As individuals engage in reflective reprocessing, they are able to formulate more complex, nuanced representations of a stimulus (e.g., allowing a stimulus to be understood in terms of multiple conflicting dimensions of evaluation).

In short, the IR Model can account for “implicit” and “explicit” attitudes with a single set of representations. The iterative reprocessing of *the same set of* stimulus information renders an increasingly conscious, reflective, and controlled evaluation, allowing the individual to better adapt to environmental challenges.<sup>6</sup> The particular level of reprocessing elicited in any particular instance of evaluation may depend on a variety of situational variables, creating the possibility for considerable

---

5. It is because evaluations are shaped from moment to moment that dissociations among attitudes measures (e.g., differences among implicit measures or between implicit and explicit measures) can be easily observed. Specifically, different attitude measures are sensitive to evaluations at different stages of evaluative processing.

6. Again, although we propose that a single set of representations can account for much of the data associated with attitudes and evaluations, we acknowledge that additional, more somewhat “symbolic,” representations may also exist. However, even if these representations exist within a separate representational store, they likely depend on the associative systems proposed in the IR Model to operate functionally.

cross-situational variability in evaluations despite the constraints on evaluation imposed by the attitude.

#### HOW MANY PROCESSES/SYSTEMS?

An important source of debate in the field of attitude research, and in psychology more generally, concerns the number of processes or systems that underlie evaluation. Although the names for the systems and explanatory domains of the models (e.g., structure, function, persuasion, etc.) often differ, most models propose two sets of processes or two systems. The first system operates relatively automatically and effortlessly, and the second system requires more cognitive attention or effort (Chaiken, 1980; Gilbert, Pelham, & Krull, 1988; Fazio, 1990; Greenwald & Banaji, 1995; Petty & Cacioppo, 1984). Further, automatic processes are thought to operate on associative information, whereas more reflective processes are thought to operate on propositional, symbolic, or linguistic representations (Gawronski & Bodenhausen, *this issue*; Lieberman, 2007).<sup>7</sup> In some models, these processes are proposed to operate relatively independently (Rydell & McConnell, 2006), whereas other models allow for interactions between systems (Gawronski & Bodenhausen, *this issue*, Strack & Deutsch, 2004, Lieberman, Gaunt, Gilbert, & Trope, 2002),

The IR Model proposes a single evaluative system in which there are many functionally discrete computational processes. This model differs from previous models in several important ways. First, by operating within a single evaluative system, these processes are free to interact in a biologically constrained fashion to resolve any current evaluative puzzle. Second, by decomposing the evaluative system into numerous computational processes, our model aims to capture the complexity of the evaluative system explicated by recent work in cognitive neuroscience. Evidence suggests that there are more than two processes involved in evaluation, and that different components within the system are performing unique computations on evaluative information. Third, the IR Model is explicit about the hierarchical organization of the components within the system. Thus, although some processes may be thought of as

---

7. It is important to note the current work in computational cognitive neuroscience suggests that this proposed fundamental and qualitative distinction between associative and propositional processing/representation may be somewhat of a false dichotomy when viewing the brain in a hierarchical fashion (Rougier, Noelle, Braver, Cohen, & O'Reilly (2005).

relatively more automatic than others, this may be more accurately described as a continuum in which some “automatic” processes are more automatic than others, and some “reflective” processes require more deliberation than others (see Cunningham & Johnson, 2007). This hierarchical approach, which views reflection as a matter of degree, is consistent with contemporary characterizations of prefrontal cortical function (Bunge & Zelazo, 2006).

Lastly and most importantly, the brain systems thought to underlie proposed dual-system models do not often operate in an all-or-none fashion. Amygdala activation, implicated in automatic evaluative processing, does not necessitate activation in the OFC or ACC (areas that are typically proposed to be part of an automatic system; Lieberman et al., 2002; Lieberman, 2007). In contrast, the IR Model seeks to understand the particular and unique computations of each region and how they may dynamically interact and be integrated in the service of rendering functional evaluations. Importantly, the precise combination and order of component processes operating in a given evaluative cycle has important and distinct implications for the final evaluation. Evidence of this claim is found in the unique constellation of evaluative deficits among patients with artificially or permanently impaired components of their evaluative systems (e.g., Damasio, 1994).

#### NATURE OF AMBIVALENCE

Adaptive functioning depends on the ability to disambiguate complex evaluative objects. How human beings process and resolve ambivalence, typically defined as a simultaneous activation of positive *and* negative associations to the same stimulus, is a critical question in the study of attitudes and evaluation. While ambivalent attitudes represent a challenge when an individual is forced to direct behavior in a binary approach/avoid fashion, they also allow for the representation of complexity and behavioral flexibility. Ambivalent attitudes may involve a lifetime of contemplation and exploration, and without resolution can provoke states of aroused inaction. Often, the root of ambivalence comes from holding two conflicting, but strong attitudes—such as when one with liberal values has to decide whether or not to give the Ku Klux Klan a permit to hold a rally. Generally, one’s egalitarian attitude does not have to be reconciled with one’s attitude toward free speech; however in this instance, the two must be resolved before a course of action can be determined. According to the IR Model, just as the concepts of attitudes and evaluation need to be carefully decomposed, so does ambivalence. Indeed, although a *subjective* sense of

ambivalence may only occur with higher levels of evaluative processing, *objective* ambivalence, which is defined as the existence of positive and negative responses to the same stimulus—which are at least the precursors of ambivalence—can exist throughout the evaluative system.

At the simplest level, then, associations with differing evaluative consequences are a precondition for ambivalence.<sup>8</sup> Within our connectionist framework, these representations exist in the connection *weights* within the network. A pattern of weights that contains both positive *and* negative information allows for the *possibility* of ambivalence and can be considered an *ambivalent attitude*. However, this potential for ambivalence is not the state of ambivalence, just as an attitude is not the same as an evaluation.

In most cases, contexts provide a solution to the problem of ambivalence. As noted by O'Reilly and Munakata (2000), the processing of context directly biases the patterns of activity, which, in this case, may result in different evaluations. Through learned conditional rules (if  $X_1$  then  $Y_1$ ; if  $X_2$  then  $Y_2$ ), likely processed in the OFC, ambivalence can be resolved by foregrounding the evaluative dimension that is more relevant to the context and backgrounding the less relevant evaluative dimension. More concretely, consider the evaluation of a cheesecake. The on-again/ off-again dieter has the potential to evaluate the dessert as a tasty treat, or as a plate full of fat. Construing the cheesecake in either of these two ways will result in vastly different evaluations. However, the possession of a "treat yourself" goal (after the latest rejection letter) or "diet" goal (after the latest trip to the gym) helps to solve the ambivalence—the context determines the current evaluation and no further processing is necessary. Interestingly, although there is large *potential* for conscious ambivalence in such cases, little ambivalence is likely felt to the extent that conflict is quickly resolved.

Subjective experiences of ambivalence arise in situations when competing stimulus representations cannot easily be resolved by foregrounding some aspect(s) of the situation. This may occur because there is no clear contextual affordance to help resolve it (as in a novel situation), or because both positive and negative elements of an attitude are equally relevant to the judgment in a given situation, as when the person mentioned above is simultaneously dieting and hungry. The cheesecake satisfies hunger, but impairs the diet. To the extent that these goals are si-

---

8. We acknowledge that ambivalence can also occur during conflicts between a network activation, based on attitudes, and network activation, based on the situation. Although such cases are important, this conflict is not a function of the attitudes, and as such is not discussed in this article.

multaneously active, behavioral choices are rendered difficult because of opposing response tendencies. Thus, although the potential and precursors for ambivalence exist at *all* stages of evaluative processing, this is the first time that one is likely to actually *experience* a subjective feeling of ambivalence. Thus, in this framework, reflective processing (mediated by the PFC) is a necessary component for the experience of ambivalence. Moreover, ambivalence is likely enhanced when behavioral options are constrained to a binary behavioral output (e.g., eating the cheesecake or waiting until the next meal in several hours).

Although the cheesecake example is relatively trivial, inaction may not be an option at times and one must resolve the conflict. We propose that this is one of the important features of the lateral PFC, and its role in the foregrounding of information. By deliberately weighting some information more than others, individuals can form a more integrated evaluation through second-order factor analysis and use it to direct behavior. Unlike inhibition, which drives inconvenient information out of mind, these integrated evaluations yield more complex activations, and may represent and account for inconsistencies. For example, the statement, "I know that the Implicit Association Test has problems, but it works really well to answer my particular question" does not hide the negative, but weighs the positives against the negatives. The evaluations that result from this type of processing are similar to what Petty and Cacioppo (1986) call an "elaborated attitude" which is known to be relatively stable and resistant to change. In this sense, these evaluations can be thought of as "resolved ambivalence."

## SUMMARY AND CONCLUSION

In determining how neuronal function gives rise to complex patterns of thought, feeling, and action, neuroscientists have paid increasing attention to the role of distributed processing, involving multiple brain regions. According to this view, information is processed in a dynamic fashion through iterative loops that progressively recruit additional brain regions. In this article, we have outlined a simplified version of an iterative circuit that might be involved in the processing and reprocessing of evaluative responses, allowing for the shaping of evaluations that are consistent with context and perceiver goals. Central to the IR Model is the hypothesis that evaluations reflect a dynamic interaction between multiple limbic and cortical structures. Automatic evaluations (which occur during earlier iterations) are linked more to limbic processing, whereas reflective evaluations additionally recruit multiple PFC regions. In addition, we have outlined a factor-analytic and computa-

tional metaphor for the types of algorithms that underlie different stages of evaluative processing. We suggest that a deeper understanding of the actual function of the evaluative system and the neural processes therein has the potential to shed light on a number of long-standing issues in the attitude literature.

## REFERENCES

- Adolphs, R. (2003). Cognitive neuroscience of human social behaviour. *Nature Reviews Neuroscience*, *4*, 165–178.
- Armony, J. L., & Dolan, R. J. (2002). Modulation of spatial attention by fear-conditioned stimuli: An event-related fMRI study. *Neuropsychologia*, *40*, 807–826.
- Bechara, A. (2004). The role of emotion in decision-making: Evidence from neurological patients with orbitofrontal damage. *Brain and Cognition*, *55*, 30–40.
- Beer, J. S., Heery, E. A., Keltner, D., Scabini, D., & Knight, R. T. (2003). The regulatory function of self-conscious emotion: Insights from patients with orbitofrontal damage. *Journal of Personality and Social Psychology*, *85*, 589–593.
- Blair, R. J. R. (2004). The roles of orbital frontal cortex in the modulation of antisocial behavior. *Brain and Cognition*, *55*, 198–208.
- Bunge, S., & Zelazo, P. D. (2006). A brain-based account of the development of rule use in childhood. *Current Directions in Psychological Science*, *15*, 118–121.
- Bush, G., Luu, P., & Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences*, *4*, 215–222.
- Cacioppo, J. T., & Berntson, G. G. (1994). Relationship between attitudes and evaluative space: A critical review, with emphasis on the separability of positive and negative substrates. *Psychological Bulletin*, *115*, 401–423.
- Cacioppo, J. T., Berntson, G. G., Sheridan, J. F., & McClintock, M. K. (2000). Multi-level integrative analyses of human behavior: Social neuroscience and the complementing nature of social and biological approaches. *Psychological Bulletin*, *126*, 829–843.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*, 747–749.
- Carver, C. S., & Scheier, M. F. (2001). *On the self-regulation of behaviour*. Cambridge, UK: Cambridge University Press.
- Chaiken S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, *39*, 752–766.
- Chaiken S., & Trope, Y. (1999). *Dual-process theories in social psychology*. New York: Guilford Press.

- Conrey, F.R., & Smith, E.R. (2007). Attitude representation: Attitudes as patterns in a distributed connectionist representational system. *Social Cognition, 25*(5), 718–735.
- Critchley, H. D., Wiens, S., Rotshtein, P., Öhman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature Neuroscience, 7*, 189–195.
- Crone, E. A., & von der Molen, M. W. (2004). Developmental changes in real life decision making: Performance on the gambling task previously shown to depend on the ventromedial prefrontal cortex. *Developmental Neuropsychology, 25*, 251–279.
- Cunningham, W. A., & Johnson, M. K. (2007). Attitudes and evaluation: Toward a component process framework. In E. Harmon-Jones & P. Winkielman (Eds.), *Social neuroscience: Integrating biological and psychological explanations of social behavior* (pp. 227–245). New York: Guilford Press.
- Cunningham, W. A., Johnson, M. K., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2003). Neural components of social evaluation. *Journal of Personality and Social Psychology, 85*, 639–649.
- Cunningham, W. A., Johnson, M. K., Raye, C. L., Gatenby, J. C., Gore, J. C., & Banaji, M. R. (2004). Separable neural components in the processing of Black and White faces. *Psychological Science, 15*, 806–813.
- Cunningham, W. A., Nezlek, J. B., & Banaji, M. R. (2004). Implicit and explicit ethnocentrism: Revisiting the ideologies of prejudice. *Personality & Social Psychology Bulletin, 30*, 1332–1346.
- Cunningham, W. A., Packer, D. J., Kesek, A., & Van Bavel, J. J. (in press). Implicit measurement of attitudes: A physiological approach. In R. E. Petty, R. H. Fazio, & P. Brinol (Eds.), *Insights from the new implicit measures*.
- Cunningham, W. A., Raye, C. L., & Johnson, M. K. (2004). Implicit and explicit evaluation: fMRI correlates of valence, emotional intensity, and control in the processing of attitudes. *Journal of Cognitive Neuroscience, 16*, 1717–1729.
- Cunningham, W. A., Raye, C. L., & Johnson, M. K. (2005). Neural correlates of evaluation associated with promotion and prevention regulatory focus. *Cognitive, Affective, & Behavioral Neuroscience, 5*, 202–211.
- Cunningham, W. A., Van Bavel, J. J., & Johnsen, I. R. (2007). Affective flexibility: Evaluative processing goals shape amygdala activity. *Psychological Science*.
- Cunningham, W. A., & Zelazo, P. D. (2007). Attitudes and evaluations: A social cognitive neuroscience perspective. *Trends in Cognitive Sciences, 11*, 97–104.
- Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Putnam.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London, B, Biological Sciences, 352*, 1413–1420.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5–18.

- Eagly, A. H., & Chaiken, S. (2007). The advantages of an inclusive definition of attitude. *Social Cognition, 25*(5), 582–602.
- Fazio, R. H. (1990). Multiple processes by which attitudes guide behavior: The MODE model as an integrative framework. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 23, pp. 75–109). New York: Academic Press.
- Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition, 25*(5), 603–637.
- 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.
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- 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.
- Gilbert, D. T., Pelham, B. W., & Krull, D. S. (1988). On cognitive business: When person perceivers meet persons perceived. *Journal of Personality and Social Psychology, 54*, 733–740.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464–1480.
- LeDoux, J. E. (1996). *The emotional brain*. New York: Simon & Schuster.
- Lieberman, M. D. (2007). The X- and C-systems: The neural basis of automatic and controlled social cognition. In E. Harmon-Jones & P. Winkelman (Eds.), *Fundamentals of social neuroscience* (pp. 290–315). New York: Guilford Press.
- Lieberman, M. D., Gaunt, R., Gilbert, D. T., & Trope, Y. (2002). Reflection and reflexion: A social cognitive neuroscience approach to attributional inference. *Advances in Experimental Social Psychology, 34*, 199–249.
- Luria, A. R. (1966). *Higher cortical functions in man* (2nd ed.). New York: Basic Books. (Original work published in 1962).
- MacDonald, A. W., Cohen, J. D., Stenger, V. A., & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science, 288*, 1835–1838.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. London: Freeman & Company.
- Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology, General, 134*, 565–584.
- Nosek, B. A., Banaji, M., & Greenwald, A. G. (2002). Harvesting implicit group attitudes and beliefs from a demonstration website. *Group Dynamics: Theory, Research and Practice, 6*, 101–115.
- Ochsner, K. N., Bunge, S. A., Gross, J. J., & Gabrieli, J. D. E. (2002). Rethinking

- feelings: An fMRI study of the cognitive regulation of emotion. *Journal of Cognitive Neuroscience*, 14, 1215–1229.
- Ochsner, K. N., & Lieberman, M. D. (2001). The emergence of social neuroscience. *American Psychologist*, 56, 717–734.
- Ochsner, K. N., Ray, R. D., Cooper, J. C., Robertson, E. R., Chopra, S., Gabrieli, J. D. E., & Gross, J. J. (2004). For better or for worse: Neural systems supporting the cognitive down- and up-regulation of negative emotion. *NeuroImage*, 23, 483–499.
- Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review*, 108, 483–522.
- O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- Osgood, C. E., Suci, G. C., & Tannenbaum, P. H. (1967). *The measurement of meaning*. Urbana, IL: University of Illinois Press.
- Oya, H., Kawasaki, H., Howard, M. A., & Adolphs, R. (2002). Electrophysiological responses in the human amygdala discriminate emotion categories of complex visual stimuli. *Journal of Neuroscience*, 22, 9502–9512.
- Packer, D. J., Kesek, A., & Cunningham, W. A. (in press). Self-regulation and evaluative processing. In A. Todorov, S. Fiske, & D. Prentice (Eds.), *Social Neuroscience: Toward understanding the underpinnings of the social mind*.
- Payne, K. B. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, 81, 181–192.
- Payne, K. B., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, 89, 277–293.
- Petty, R. E., & Cacioppo, J. T. (1984). The effects of involvement on response to argument quantity and quality: Central and peripheral routes to persuasion. *Journal of Personality and Social Psychology*, 46, 69–81.
- Petty, R. E., Briñol, P., & DeMarree, K. G. (2007). The Meta-Cognitive Model (MCM) of attitudes: Implications for attitude measurement, change, and strength. *Social Cognition*, 25(5), 657–686.
- Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 123–205). New York: Academic Press.
- Phelps, E. A. (2006). Emotion and cognition: Insights from studies of the human amygdala. *Annual Review of Psychology*, 24, 27–53.
- Phelps, E. A., O'Connor, K. J., Gatenby, J. C., Grillon, C., Gore, J. C., & Davis, M. (2001). Activation of the human amygdala to a cognitive representation of fear. *Nature Neuroscience*, 4, 437–441.
- Priester, J. M., & Petty, R. E. (1996). The gradual threshold model of ambivalence: Relating the positive and negative bases of attitudes to subjective ambivalence. *Journal of Personality and Social Psychology*, 71, 431–449.

- Ridderinkhof, K. R., Ullsperger, M., Crone, E. A., & Nieuwenhuis, S. (2004). The role of the medial frontal cortex in cognitive control. *Science*, *306*, 443–447.
- Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex*, *10*, 284–294.
- Rolls, E. T., Hornak, J., Wade, D., & McGrath, J. (1994). Emotion-related learning in patients with social and emotional changes associated with frontal lobe damage. *Journal of Neurology, Neurosurgery & Psychiatry*, *57*, 1518–1524.
- Rougier, N.P., Noelle, D., Braver, T.S., Cohen, J.D., & O'Reilly, R.C. (2005). Prefrontal cortex and the flexibility of cognitive control: Rules without symbols. *Proceedings of the National Academy of Sciences*, *102*, 7338–7343.
- 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.
- Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition*, *25*(5), 638–656.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London, B*, *298*, 199–209.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, *8*, 220–247.
- Stuss, D. T., & Benson, D. F. (1986). *The frontal lobes*. New York: Raven Press.
- Whalen, P. J. (1998). Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. *Current Directions in Psychological Science*, *7*, 177–188.
- Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In L. Martin & A. Tesser (Eds.), *The construction of social judgment* (pp. 37–65). Hillsdale, NJ: Erlbaum.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, *107*, 101–126.
- Zelazo, P. D. (2004). The development of conscious control in childhood. *Trends in Cognitive Science*, *8*, 12–17.
- Zelazo, P. D., & Cunningham, W. (2007). Executive function: Mechanisms underlying emotion regulation. In J. Gross (Ed.), *Handbook of emotion regulation* (pp. 135–158). New York: Guilford Press.