CHAPTER 1

A History of Implicit Social Cognition

Where Is It Coming From? Where Is It Now?
Where Is It Going?

B. Keith Payne and Bertram Gawronski

Within the space of two decades, virtually every intellectual question in social psychology, and many outside of it, has been shaped by the theories and methods of *implicit social cognition*. Many of those questions are pondered in this volume, involving the role of automatic/implicit/unconscious processes in attitudes (Petty & Briñol, Chapter 18), social judgment and decision making (Bodenhausen & Todd, Chapter 15), goal pursuit (Ferguson & Porter, Chapter 17), prejudice and stereotyping (Amodio & Mendoza, Chapter 19; Trawalter & Shapiro, Chapter 20), self-concepts and self-esteem (Schnabel & Asendorpf, Chapter 22; Zeigler-Hill & Jordan, Chapter 21), social cognitive development (Olson & Dunham, Chapter 13), romantic relationships (Baldwin, Lydon, McClure, & Etchison, Chapter 23), and social justice (Payne & Cameron, Chapter 24). Expanding beyond the traditional boundaries of social psychology, the basic ideas of implicit social cognition have also had an impact in many applied areas, including health psychology (Wiers et al., Chapter 25), clinical psychology (Teachman, Cody, & Clerkin, Chapter 26), forensic psychology (Snowden & Gray, Chapter 27), consumer psychology (Perkins & Forehand, Chapter 28), and political psychology (Nosek, Graham, & Hawkins, Chapter 29).

In every topic of study, implicit social cognition is concerned with automatic/implicit/unconscious processes underlying judgments and social behavior. An indispensable part of this endeavor is the use of a new class of indirect measurement procedures such as the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998) and different kinds of sequential priming tasks (e.g., Fazio, Jackson, Dunton, & Williams, 1995; Payne, Cheng, Govorun, & Stewart, 2005; Wittenbrink, Judd, & Park, 1997), which play a crucial role in the chapters of this book. Thumbing through its pages makes it clear how implicit social cognition has brought new insights, and also new controversies, wherever it has led. In fact, implicit social cognition has grown at such an incredible rate over the past years that it seems almost impossible to keep track of the rapidly progressing developments in this area. The purpose of this handbook is to provide a comprehensive overview of the current state of the field, including key findings, current directions, and emerging themes in the area of implicit social cognition.

In this introduction, we aim to provide a context for the chapters that follow by highlighting some of the themes that keep surfacing in the field, tracing them to their historical roots, and identifying

emerging themes that may guide future research. We hope that our historical synopsis will not only provide an orientation for the in-depth discussions of particular topics in the following chapters but also circumscribe the characteristics that define implicit social cognition as a particular way of studying, understanding, and explaining human behavior. In addition, we hope that our introduction helps illuminate the historical roots of previous and ongoing debates, which seems valuable for critical appraisals of theoretical interpretations in implicit social cognition.

WHERE IS IT COMING FROM?

A history of implicit social cognition could easily enough start with Freud and the psychoanalytic unconscious; or with Augustine and Aquinas, who had a lot to say about the limits of introspection in knowing thyself; or even with Plato and Aristotle, who commented extensively on consciousness and intentional behavior. Our aim in this introduction is more modest. It is to take a selective look at roughly the last two decades in which implicit social cognition has comprised a recognizable enterprise in its own right, establishing itself as one of the most influential approaches in social psychology. This process began as researchers adapted ideas and methods from cognitive psychology to answer social psychological questions. As we shall see, many of the lively debates in implicit social cognition can be traced to which particular tradition of cognitive psychology was appropriated.

Two Roots of Implicit Social Cognition

Reports on implicit social cognition include a lot of hyphens and slashes in their key terms. Phrases like *automatic/implicit/unconscious processing* and *controlled/explicit/conscious processing* abound. Aside from being less than precise, such phrases reveal a tension at the heart of the discipline. One theme that recurs in the field, and in this book, is an uneasy relationship between notions of automaticity and unconsciousness. When these terms are thrown together, they sometimes leave readers with only a vague idea of what is being studied. However, a sharper picture emerges when we take a step back and distinguish two intellectual traditions that gave rise to separate terminologies. One tradition is found in research on selective attention. The other is found in research on implicit memory.

Roots of Automatic and Controlled Processing in Attention Research

Some of the seminal articles that can be subsumed under the umbrella term *implicit social cognition* emphasize the differences between automatic and controlled cognition, with little mention of the unconscious. These reports took as their point of departure the cognitive psychology work of Shiffrin and Schneider (1977) and Posner and Snyder (1975), among others. This work developed from a tradition of research on selective attention and short-term memory (e.g., Broadbent, 1971; Treisman, 1969). The key idea in these articles was that information processing could be divided into controlled and automatic modes. Controlled processing was defined as demanding attention, limited in capacity, and voluntarily initiated and altered. Automatic processing was defined as needing little attention, unlimited in capacity, and difficult to suppress voluntarily. These criteria will, of course, sound very familiar to social cognition researchers because they are the forerunners, but not identical to, Bargh's (1994) "four horsemen" of automaticity (awareness, efficiency, intention, control). Although in Bargh's scheme conscious awareness is the first feature, the earlier cognitive theories hardly mentioned consciousness.

Fazio and colleagues' seminal work showing that attitudes can be automatically activated built on these cognitive theories of automatic and controlled processing (Fazio, Sanbonmatsu, Powell, & Kardes, 1986; see also Dovidio, Evans, & Tyler, 1986; Gaertner & McLaughlin, 1983). Here too there was little concern with consciousness. The term *conscious* appears only once in Fazio and colleagues' (1986) article and the term *aware* appears twice. For Fazio and colleagues, "The key feature of such automatic activation, then, is inescapability" (p. 229). In the work of Shiffrin and Schneider (1977) and Posner and Snyder (1975), the degree of learning was critically important for automaticity. Well-learned items were detected or retrieved from memory automatically, whereas poorly learned items required cognitive effort to search for them. Based on these assumptions, Fazio and colleagues distinguished between well-learned (i.e., strong) attitudes that should be activated automatically and poorly learned (i.e., weak) attitudes that should not.

The connection between the strength of attitudes and their ability to be automatically activated set the stage for the use of sequential priming techniques to measure attitudes without asking subjects to report them (Fazio et al., 1995; see also

Wentura & Degner, Chapter 6, this volume). This novel approach to indirect measurement would soon become important across many areas of psychology and ultimately a central characteristic of implicit social cognition. Initially, the usefulness of this approach was most obvious for studying racial attitudes (see Amodio & Mendoza, Chapter 19, and Trawalter & Shapiro, Chapter 20, this volume). Racial attitudes had always presented challenges for researchers because self-presentation motives meant many subjects would not honestly report their attitudes. Techniques such as the bogus pipeline showed that subjects typically hold more prejudiced attitudes than they are willing to report under ordinary circumstances (Jones & Sigall, 1971). However, such deceptive procedures are ethically controversial and practically cumbersome. Priming methods seemed to offer the potential for a bona fide pipeline to reach respondents' true attitudes (Fazio et al., 1995).

The challenge of racial attitudes was taken up by both Devine (1989) and Fazio and colleagues (1995). Both studies were informed by cognitive theories of automatic and controlled processing (Posner & Snyder, 1975; Shiffrin & Schneider, 1977) and priming techniques (Neely, 1977); both had little to say about conscious awareness of attitudes; and both emphasized the idea that well-learned associations should be activated automatically, but weakly learned associations require cognitive effort to be retrieved. For Fazio, the implication was that individuals with weak or neutral racial attitudes simply would not show any priming effects, whereas those with strongly negative attitudes or strongly positive attitudes would show corresponding priming effects. However, for Devine the important distinction was between the *knowledge* of a social stereotype versus the personal *endorsement* of the stereotype. Because everyone in a culture learns stereotypes in the same way they learn about other categories, such as birds or foods, knowledge of stereotypes should be well learned for virtually everyone. Stereotypes should, therefore, be automatically activated for everyone. Personally endorsed beliefs, on the other hand, should vary with one's values and motivations. Thus, in Devine's approach, inescapable stereotype activation was universal; it was endorsed beliefs that distinguished prejudiced from unprejudiced individuals.

Here we see a second theme that reappears in debates throughout implicit social cognition research: Are the results of cognitive measurements, such as the ones revealed by sequential priming tasks, best considered a reflection of the person's inner attitudes and beliefs, or are they instead the products of the cultural environment? This debate is complex and multifaceted, and it recurred in different variants. Examples include the early controversy between Devine (1989) and Fazio and colleagues (1995), debates about the generality versus variability of automatic attitude activation (Bargh, Chaiken, Govender, & Pratto, 1992; Chaiken & Bargh, 1993; Fazio, 1993), and recent controversies about the role of personal versus extrapersonal associations in the IAT (Gawronski, Peters, & LeBel, 2008; Nosek & Hansen, 2008a, 2008b; Olson & Fazio, 2004; Olson, Fazio, & Han, 2009). We consider three different aspects of the debate in turn. First, the debate has sometimes been understood as questioning whether automatic attitudes are "real." Both Fazio and Devine showed that automatically activated stereotypes or attitudes were associated with biased impressions and behaviors. Insofar as the reality of an attitude is a question about whether it has consequences for thought and behavior, all sides of the debate agree that the automatic responses are real.

A second aspect of this debate is whether there are meaningful individual differences in automatic responses. If automatic responses were like language and virtually every healthy member of a culture possessed them, variability in sequential priming scores would simply reflect measurement error. A great deal of subsequent research showed that individual differences in priming and other cognitive tasks are indeed meaningful and predictive of behavior (Perugini, Richetin, & Zogmaister, Chapter 14, this volume).

A thornier aspect of this debate is whether people who show negative automatic associations with ethnic groups should be considered prejudiced (Arkes & Tetlock, 2004; Banaji, Nosek, & Greenwald, 2004). This is a philosophical and moral issue more than an empirical one. Not surprisingly, it remains a topic of hot debate (see Payne & Cameron, Chapter 24, this volume). The roots of the debate can be traced as far as concepts of virtue in ancient Greek philosophy. Plato argued that virtue is about developing the right inclinations. The virtuous person has good inclinations, whereas a wicked person has evil inclinations. However, Aristotle argued that both good and bad people sometimes have wicked impulses, and that virtue consists in regulating and channeling them well. Contemporary philosophers have begun integrating insights from implicit social cognition with theories of ethics, creating an exciting exchange of ideas across disciplines (e.g., Appiah, 2008; Doris, 2002; Kennett & Fine, 2009).

Roots of Unconscious and Conscious Processing in Implicit Memory Research

The studies we have discussed so far focused on the distinction between automatic and controlled processing, emphasizing notions of inescapability rather than unconsciousness. However, in the following years, the automatic/controlled distinction often gave way to the dichotomy between explicit and implicit processes, interpreted as synonyms for the terms *conscious* and *unconscious*. This shift was ignited in part by Greenwald and Banaji's (1995) influential review, which built on a different research tradition of cognitive psychology. Rather than building on theories of attention and short-term memory, they drew on research on implicit memory (Banaji, 2001). Although measured in a variety of ways, implicit memory has been defined fairly consistently as influences of past experience on later performance, in the absence of conscious memory for the earlier experience (Jacoby & Dallas, 1981; Schacter, 1987). Amnesic patients, for example, tend to perform at chance when asked to recognize which words they have previously studied. When asked, however, to guess how to complete word fragments, they perform more accurately for studied than for unstudied words (Warrington & Weiskrantz, 1968). The effect of prior study on later performance, despite the inability to consciously remember it, is taken to indicate implicit memory.

This conceptualization strongly shaped Greenwald and Banaji's (1995) definition of implicit attitudes as "introspectively unidentified (or inaccurately identified) traces of past experience that mediate favorable or unfavorable feeling, thought, or action toward social objects" (p. 8). Just as implicit memory was defined as consequences of past experience in the absence of conscious awareness for the experience, implicit attitudes were defined as traces of past experience in the absence of conscious awareness for the experience. Implicitness was identified with unconsciousness. In contrast to earlier reports emphasizing the notion of automaticity in the sense of inescapability (Fazio et al., 1986), variants of the words *conscious* and *awareness* are mentioned nearly a hundred times in Greenwald and Banaji's paper.

Although the parallel between concepts of implicit memory and implicit attitudes is close, it is not complete. There was a subtle shift between the two concepts that has important consequences for implicit social cognition. In the case of implicit memory, it is the experience of some past event that is not consciously available. In memory stud-

ies, the "event" is often the presentation of a particular word or a picture in a learning task. The experimenter has perfect control over whether a particular stimulus was presented and, therefore, knows with certainty whether the subject experienced a particular learning episode. Moreover, the definition of implicit memory focuses on performance, which is an operational definition. From this perspective, implicit memory is based on the effects of past experiences, controlled by the experimenter, on later performance, which is directly observable. For implicit attitudes, in contrast, the definition refers to *traces* of past experience, which *mediate* later responses. This definition is decidedly more mentalistic than for implicit memory. In fact, implicit social cognition researchers usually do not control conscious awareness of past experiences giving rise to attitudes, stereotypes, and other such traces. Unlike implicit memory tests, the measures used in implicit social cognition research do not test for awareness of the formative experiences. Moreover, although Greenwald and Banaji (1995) clearly referred to unawareness of the experiences giving rise to attitudes, subsequent writers have sometimes conflated awareness of the source with awareness of the attitude itself. This conflation has led to the fallacy that if one assesses a construct with a measure that does not presuppose conscious introspection, then the assessed construct must be introspectively inaccessible. The available evidence clearly speaks against this assumption (Gawronski, Hofmann, & Wilbur, 2006), but claims of unconsciousness remain common. The degree to which the constructs measured with cognitive tests may be unconscious and how to empirically confirm such claims remains a topic of debate. This debate, in turn, has inspired novel ways of thinking about consciousness and introspection (Hofmann & Wilson, Chapter 11).

What Is "Implicit" about Implicit Social Cognition?

Since Greenwald and Banaji (1995) coined the term *implicit social cognition*, the two roots of implicit social cognition have contributed to lively debates about the proper use of terminology, in particular with regard to the term *implicit*. Some researchers have used the term to describe a particular characteristic of measurement procedures, namely measures that provide indicators of psychological attributes (e.g., attitudes) without having to ask participants to verbally report the desired information (e.g., Fazio & Olson, 2003). Yet

other researchers have used *implicit* to describe the constructs assessed by a particular class of measurement procedures, namely constructs assessed by tasks that do not require conscious introspection and, therefore, might reflect psychological attributes that are introspectively inaccessible (e.g., Banaji, 2001).

Because such terminological issues are a matter of semantics and linguistic conventions, they cannot be resolved empirically. At the same time, different use of the same terms often produces conceptual confusion, which can have disruptive effects on cumulative research progress (Machado & Silva, 2007). To overcome these problems, researchers have developed detailed taxonomies of the various properties of measurement procedures (De Houwer & Moors, Chapter 10, this volume) and mental representations (Carlston, Chapter 3, this volume). The most comprehensive conceptual analysis has been provided by De Houwer and colleagues, who suggested using the terms *direct* and *indirect* to describe features of measurement procedures and the terms *implicit* and *explicit* to describe features of the psychological attributes that are assessed by different measurement procedures (e.g., De Houwer, Teige-Mocigemba, Spruyt, & Moors, 2009). With regard to the term *implicit*, De Houwer further suggested using this description in the meaning of the term *automatic*, such that it may describe the unconscious, efficient, unintentional, or uncontrollable nature of the assessed constructs (Moors, Spruyt, & De Houwer, Chapter 2, this volume).

Even though terminological systems like the one proposed by De Houwer are a matter of convention, the normative request implied by these taxonomies is the call for researchers to be precise about what they mean when they use the term *implicit*. These norms encourage researchers to specify in which particular sense a process is automatic (De Houwer et al., 2009) and what exactly remains outside of conscious awareness (e.g., past experiences that give rise to attitudes vs. attitudes per se; see Gawronski et al., 2006). Detailed taxonomies and clarity in "defining our terms" will not themselves settle empirical questions, but they can strongly facilitate research progress by preventing conceptual misunderstandings.[1]

A NEW INDUSTRY OF RESEARCH

Greenwald and Banaji (1995) ended their review with a call for the refinement of individual difference measures of implicit social cognition. They predicted that "when such measures do become available, there should follow the rapid development of a new industry of research on implicit cognitive aspects of personality and social behavior" (p. 20). Their prediction has not been disappointed. With the development of the IAT (Greenwald et al., 1998), implicit social cognition research seemed to hit a tipping point. Although sequential priming tasks had been available for several years (e.g., Dovidio et al., 1986; Fazio et al., 1986; Gaertner & McLaughlin, 1983), researchers now had a task that made implicit social cognition research much easier. Facilitation scores from priming procedures often suffered from low reliability and relatively small effect sizes. However, the IAT had reasonably good reliability and very large effects. In fact, even though the term *implicit social cognition* was originally interpreted in a much broader sense (Greenwald & Banaji, 1995), it has become almost synonymous with research using indirect measurement procedures, such as the IAT and its derivates (Teige-Mocigemba, Klauer, & Sherman, Chapter 7, this volume), sequential priming tasks (Wentura & Degner, Chapter 6, this volume), and sophisticated indirect paper-and-pencil methods (Sekaquaptewa, Vargas, & von Hippel, Chapter 8, this volume). When the number of studies using indirect measurement procedures started to skyrocket around the turn of the millennium, the different roots of implicit social cognition left their own traces by inspiring two parallel, largely independent streams of research.

Two Parallel Streams of Research

The first stream had its origin in the idea of inescapability, derived from the notion of automaticity in attention research. This stream was mainly theory driven, in that empirical predictions were derived from established theories of attitude–behavior relations, most notably Fazio's motivation and opportunity as determinants (MODE) model (Fazio, 1990; Fazio & Towles-Schwen, 1999). The central assumption in the MODE model is that direct and indirect measures assess the behavioral effects of the same underlying mental representation, the only difference being the degree of control that participants have over their responses. Whereas responses on self-report measures can be easily altered, indirect measures such as the IAT and sequential priming tasks constrain participants' opportunity to control their responses. This distinction captures the theoretical core of the MODE model, which argues that automatically activated attitudes will guide judgments and

behavior if either the motivation or the opportunity to engage in deliberate processing is low. If, however, both motivation and opportunity are high, the impact of automatic attitudes on self-reports may be diluted or inhibited by deliberate processes. These assumptions imply that the predictive validity of a given measure should depend on the overlap between the processing conditions during the completion of the measure and the processing conditions of the relevant behavior. For instance, indirect measures should be better predictors of spontaneous behavior, whereas direct self-report measures should be better predictors of deliberate behavior (Perugini et al., Chapter 14, this volume).

A second, parallel stream of research had its origin in the notion of unconsciousness, as derived from cognitive research on implicit memory. In contrast to the single-representation assumption implied by the MODE model, this stream was characterized by an operational equation of measurement outcomes with distinct psychological constructs, which favored a dualism of two independent representations in memory: a conscious, explicit representation and an unconscious, implicit representation (Greenwald & Banaji, 1995). Because there was no theoretical framework available that could specify how the two kinds of representations may differentially influence judgments and behavior, empirical research in this stream was largely effect focused and method driven. The main topics of interest were whether indirect measures predict behaviors that direct measures do not predict or whether indirect measures explain variance in behavior over and above variance that can be explained by direct measures. As reviewed by Perugini and colleagues (Chapter 14, this volume), research has documented several distinct patterns by which indirect and direct measures may differentially predict behavior. For example, in some cases one measure or the other alone predicted behavior; in other cases they predicted behavior additively; and in other cases they synergistically combined to predict behavior. There is replicable evidence for each of these patterns, but this seems to be a case in which theory has not yet caught up to the data. Many models such as the MODE model and others reviewed in this chapter can accommodate these findings, but it is not clear that any current models make unambiguous a priori predictions for when each pattern should be expected. Developing such theories is an important next step we hope researchers will pursue.

Unexpected Malleability

The explosion of research using indirect measures also led to discoveries that forced a reexamination of some of the field's core assumptions. One such discovery was that the scores obtained with indirect measures often increased, decreased, or even reversed as a function of the context (Gawronski & Sritharan, Chapter 12). Spending 5 minutes imagining a "strong woman" led to weaker implicit gender stereotyping (Blair, Ma, & Lenton, 2001), and viewing photos of admired African Americans such as Martin Luther King, Jr., along with disliked white Americans such as serial killer Jeffrey Dahmer reduced implicit race bias (Dasgupta & Greenwald, 2001). Racial bias on indirect measures was even reversed by the background context. In one study, pictures of African Americans in the role of prisoners evoked negative implicit evaluations, but presenting the same individuals in the role of lawyers elicited positive implicit evaluations (Barden, Maddux, Petty, & Brewer, 2004; see also Judd, Simpson, & Carver, 2001; Wittenbrink, Judd, & Park, 2001).

The ease of shifting measurement scores challenged a dominant assumption at the time that the associations assessed by indirect measures had become automatized by a long history of learning and must, therefore, be difficult to change (e.g., Wilson, Lindsey, & Schooler, 2000; see also Olson & Dunham, Chapter 13, this volume). This assumption was based on earlier research using perceptual learning tasks, and Shiffrin and Schneider (1977) had indeed shown that information was retrieved automatically only after hundreds or thousands of trials of learning. Thus, findings that the scores of indirect measures could be swayed one way or the other by immediate changes in context presented a puzzle. The two answers that have been offered to resolve this puzzle introduced a third theme to implicit social cognition, namely the debate about whether responses on indirect measures reflect stable representations in long-term memory or fluid constructions that are generated on the spot. In the eyes of construction theorists, the high malleability of indirect measurement scores confirmed their assumption that contexts influence what information is used to construct an attitude from one moment to the next, and that these principles apply equally to direct and indirect measurement procedures (e.g., Schwarz, 2007). In fact, the very idea that indirect measurement procedures would assess rigid "things" in memory independent of the context was seen as ill founded.

In contrast to this view, representation theorists argued that malleability in measurement scores does not reflect context-dependent changes in mental representations but, rather, responses to fundamentally different objects (e.g., Fazio, 2007). According to this view, contextual information simply influences how a given object is categorized. To the degree that different categories are associated with different evaluations in long-term memory, contextual cues can lead to different responses to the same object, even though the respective category associations that underlie these responses are highly stable. As Schwarz and Bohner (2001) pointed out, this debate cannot be resolved on empirical grounds because any result predicted by one theoretical framework can be reinterpreted by the other one. Yet, inspired by this debate, researchers have started to study the learning mechanisms that lead to context-dependent or context-independent responses on indirect measures (e.g., Rydell & Gawronski, 2009). By shifting the focus from mere demonstrations of context effects to experimentally controlled learning experiences and their directly observable effects on performance, this line of research could possibly lay the groundwork for a new way of studying context effects on indirect measures that goes beyond the debate between representational and constructionist models.

Unexpected Correspondence to Self-Report Measures

If the context dependence of indirect measurement scores made researchers rapidly reexamine their assumptions, another cause for reexamination accumulated more slowly over several years. Early studies often found little or no correspondence between direct and indirect measurement scores (Rydell & McConnell, Chapter 16, this volume). This divergence reinforced the idea that indirect measures reflected unconscious representations that are inaccessible to verbal self-report. If implicit representations were unconscious, then self-reports would certainly not be expected to track them. However, as the field's data base grew, more and more studies showed surprisingly high correspondence under certain conditions (Hofmann, Gschwendner, Nosek, & Schmitt, 2005).

One of the critical factors, identified in early research under the framework of the MODE model (Fazio et al., 1995), was the motivation to control the expression of one's attitudes. For instance, in the context of racial attitudes the correlation

between directly and indirectly assessed evaluations was much higher for those who were willing to openly express prejudice compared with those who were motivated to control prejudiced reactions (e.g., Dunton & Fazio, 1997; Fazio et al., 1995; Gawronski, Geschke, & Banse, 2003; Payne, 2001). Expanding on these findings, research has shown that the spontaneity of self-reports plays a significant role in this regard, such that direct and indirect measurement scores show higher correspondence when self-reports are based on quick, intuitive gut reactions rather than deliberate analyses (e.g., Gawronski & LeBel, 2008; Hofmann, Gawronski, Gschwendner, Le, & Schmitt, 2005; Jordan, Whitfield, & Zeigler-Hill, 2007). A second critical factor was measurement error. Cunningham, Preacher, and Banaji (2001) demonstrated that when latent-variable analyses were used to correct for measurement error, the correlation between directly and indirectly assessed racial attitudes was quite substantial (see also Hofmann, Gawronski, et al., 2005). A third factor concerns whether the attitude objects evaluated in direct and indirect measures are similar or not. Conceptual correspondence (Hofmann, Gawronski, et al., 2005) refers to whether different measures assess conceptually comparable constructs; structural fit (Payne, Burkley, & Stokes, 2008) refers to whether direct and indirect measures involve similar stimuli and responses. Overall, measurement scores of direct and indirect tests show much higher correspondence when they are matched in their concepts and test structures than when they do not align. Finally, the correspondence of measurement scores depends on several aspects of the underlying representations. For instance, direct and indirect measurement scores are more strongly related when the assessed attitudes are strong and perceived as distinctive (Nosek, 2005).

The discovery of these moderators shifted the question from whether responses on direct and indirect measures are related to identifying the conditions under which they are related (Hofmann, Gschwendner, et al., 2005). This research suggested that direct and indirect measures, in fact, reveal highly similar results when the tasks are reliable, the stimulus materials are presented in similar ways, and participants base their self-reports on quick, intuitive gut reactions rather than deliberate analyses. Taken together, these findings posed a challenge to earlier claims that indirect measures assess unconscious representations (Gawronski, LeBel, & Peters, 2007). In fact, unconsciousness of a mental representation, like

the nonexistence of black swans or the absence of light in a closed refrigerator, is notoriously difficult to prove. Even when there is no evidence for their presence, it is always possible you have not looked in the right place or at the right time (Popper, 1934). But together, these findings suggested that, as long as researchers look in the right way, people seem to have much greater introspective access to their mental representations than was commonly assumed.

SECOND-GENERATION MEASURES

Most of the early research on implicit social cognition used either the IAT (Greenwald et al., 1998) or available variants of sequential priming (e.g., Fazio et al., 1995; Wittenbrink et al., 1997). However, both types of indirect measures had their problems. Whereas sequential priming tasks often suffered from low measurement reliability, the IAT involved a number of structural problems that produced method-related confounds. These issues inspired researchers to develop new measurement procedures that tried to overcome the problems of existing tasks. The result was a second generation of measurement procedures that expanded on early work on the IAT and priming tasks.

IAT-Inspired Methods

An early recognized limitation of the IAT was the fact that it provides relative assessments involving two target objects (Teige-Mocigemba et al., Chapter 7, this volume). For instance, if a person shows a preference for candy bars over apples, we cannot tell whether he or she has a particular yen for Snickers, an aversion to Granny Smiths, or some combination of both (Blanton, Jaccard, Gonzales, & Christie, 2006; Fiedler, Messner, & Bluemke, 2006). Nosek and Banaji (2001) were the first to address this issue by developing the Go/No-Go Association Task (GNAT). The GNAT includes one target object, such as apples, and attributes, such as good and bad words. In one block of trials, subjects press a key when they see apples or good words, and in the other block they press a key when they see apples or bad words. If subjects are faster and more accurate when apples are paired with good rather than bad words, it is inferred that they have a positive attitude toward apples. Related approaches to measuring attitudes toward single-target objects have been developed by simplifying the IAT's dual-target structure to a single

target, as in the Single-Category IAT (Karpinski & Steinman, 2006) and the Single-Target IAT (Wigboldus, Holland, & van Knippenberg, 2004).

Because there is only one target object, these measures avoid the problem of relative comparisons. However, these tasks still involve comparisons across separate blocks of compatible versus incompatible trials, exposing them to another set of methodological critiques (Teige-Mocigemba et al., Chapter 7, this volume). In a nutshell, these tasks assume that participants' attitudes, stereotypes, or other knowledge structures provide the source of compatibility effects. That is, candy bars and pleasant words are compatible because both are positively evaluated. However, the valence of the items may be only one source of compatibility effects. As noted by Rothermund and Wentura (2001), participants might pair the items along any dimension that is salient at the time. Hence, participants might simplify the sorting task by subjectively recoding the task using any salient heuristic. Such "salience asymmetries" could create compatibility effects on the tasks for reasons that are unrelated to the associations of interest to the researcher. This issue has been addressed by several new methods that present congruent and incongruent trials randomly in a single block rather than blocked. The first IAT-derived task that avoided its original block structures is De Houwer's (2003) Extrinsic Affective Simon Task. More recent examples include the Single-Block IAT (Teige-Mocigemba, Klauer, & Rothermund, 2008) and the Recoding-Free IAT (Rothermund, Teige-Mocigemba, Gast, & Wentura, 2009).

Priming-Inspired Methods

These methods all aimed at remedying some specific limitations of the IAT, perhaps because the IAT's popularity made it an exciting target for refining measurement techniques. However, during the same period, research was progressing to refine priming methods as well (Wentura & Degner, Chapter 6, this volume). Whereas Fazio and colleagues (1986) had used sequential priming to measure attitudes, Banaji and Hardin (1996) applied sequential priming to measure semantic relations (e.g., stereotypes), and Wittenbrink and colleagues (1997) extended priming paradigms to measure semantic and affective dimensions of meaning independently.

Priming paradigms have the considerable benefit of simplicity. Subjects typically respond to the target items with a simple judgment, making these paradigms simpler for subjects to complete and for

researchers to interpret compared with IAT-related tasks. However, the response latency scores yielded by priming tasks often involve large proportions of measurement error. One way around this problem is to require participants to respond quickly, thus shifting influences of the primes from response latencies to accuracy rates (Draine & Greenwald, 1998). Payne (2001) used this strategy to measure the relationship between racial stereotypes and perceptions of weapons. Target items were more likely to be misidentified as guns when they were preceded by black face primes than white face primes. Priming effects in accuracy were found to be more reliable than priming effects in response latencies (Payne, 2001, 2005).

Another approach to this problem is to dispense with latency and accuracy altogether and instead use judgments of ambiguous stimuli. Murphy and Zajonc (1993) showed that participants rated Chinese ideographs more favorably when they were primed with emotionally pleasant pictures and less favorably when they were primed with unpleasant pictures. Payne, Cheng, and colleagues (2005) adapted this paradigm to measure individual differences in attitudes toward the prime stimuli. Their affect misattribution procedure proved to have large effect sizes and high reliability, helping to address some of the key difficulties with sequential priming methods.

THEORETICAL DEVELOPMENTS

Parallel to the development of second-generation measures, the field of implicit social cognition has been reshaped by two major theoretical advances. One involves the emergence of generalized, domain-independent dual-process models; the other involves the development of formal process models that describe and quantify the contributions of distinct processes to performance on cognitive tasks. Together, these theoretical developments have had a significant impact on how researchers interpret the scores obtained with indirect measures.

Generalized Dual-Process Models

At the dawn of the new millennium, social psychology witnessed a remarkable shift in the generality of theorizing. Research in the 1980s and 1990s was characterized by a strong influence of dual-process theories, which divided the realm of social cognitive processes into effortless, automatic processes versus effortful, controlled processes. To

a large extent, these theories were domain specific in that they were concerned with particular phenomena in social psychology, including persuasion (e.g., Chaiken, 1980; Petty & Cacioppo, 1986), impression formation (e.g., Brewer, 1988; Fiske & Neuberg, 1990), and dispositional attribution (Gilbert, 1989; Trope, 1986). In 2000, Smith and DeCoster published an influential review article in which they reconceptualized the domain-specific processes proposed by earlier theories in terms of a general set of processes underlying a variety of phenomena. Their distinction between associative and rule-based processes set the foundation for a theoretical reinterpretation of direct and indirect measures as reflecting the outcomes of two qualitatively distinct processes (e.g., Gawronski & Bodenhausen, 2006; Rydell & McConnell, 2006; Strack & Deutsch, 2004; see also Deutsch & Strack, Chapter 4, this volume). The important difference from earlier theorizing is that these models emphasize the principles by which processes operate (e.g., associative vs. rule based) rather than the conditions under which they operate (e.g., unconscious, efficient, unintentional, uncontrollable). These operating conditions then become empirical questions to be tested.

The most influential model in this regard is Strack and Deutsch's (2004) reflective–impulsive model, which distinguishes between associative and propositional processes. Associative processes involve the activation of associations in memory, which is guided by the principles of similarity and spatiotemporal contiguity. Propositional processes, in contrast, are conceptualized as the validation of the information implied by activated associations, which depends on syllogistic principles of logical consistency (Gawronski & Strack, 2004). The central difference between the two kinds of processes lies in the role of subjective truth or accuracy. Whereas the activation of associations is assumed to occur independently of whether these associations are regarded as accurate or inaccurate, propositional processes are inherently concerned with the validation of activated information. This distinction has been related to the outcomes of direct and indirect measures. Direct self-report measures assess the subjective validity of propositional statements about some state of affairs (e.g., "How much do you agree with the statement ... "), whereas indirect measures assess the activation—that is, the momentary accessibility—of associations independent of whether they are regarded as accurate or inaccurate.

Despite their impact on the field, generalized dual-process models have also inspired some

controversy, with some authors arguing that the distinctions they draw are too narrow and others arguing that they are too broad. For example, Kruglanski, Erb, Pierro, Mannetti, and Chun (2006) argued that it is unnecessary to distinguish between processes on the basis of operating principles such as associative versus propositional processing because virtually all judgment can be interpreted as the result of a single inference process so long as researchers specify parameters, such as relevance of information, processing difficulty, and cognitive resources needed to reach a judgment. In contrast, Sherman (2006) argued that more than two processes are often needed to describe the interplay between automatic and controlled processes. Still, generalized dual-process models have advanced the field, at a minimum, by providing heuristically useful frameworks for interpreting distinctions between indirect and direct measures (Hofmann, Gawronski, et al., 2005). In addition, the models have provided theoretical links to recent developments in social neuroscience (Ito, Chapter 5, this volume), and they inspired novel ways of thinking about reflective versus impulsive determinants of social behavior (Strack & Deutsch, 2004) and the malleability of associative representations (Gawronski & Bodenhausen, 2006).

Formal Process Models of Indirect Measures

Whereas generalized dual-process models seek to explain a wide variety of findings on the basis of a broad central distinction, another new direction shaping implicit social cognition aims to get very specific about the processes underlying individual effects, tasks, or behaviors. Formal process models describe and quantify the contributions of multiple processes to performance on particular tasks and have recently been applied to indirect measures (Sherman, Klauer, & Allen, Chapter 9, this volume). Indirect measures operate on the premise that their scores are driven by underlying attitudes, stereotypes, or whatever representation is of interest to the researcher. However, a great deal of research suggests that the reality is more complicated than that. Many variables have been shown to affect task performance besides the mental contents they were meant to measure. For example, Payne (2001) observed greater stereotype bias in weapon misidentifications when participants responded under time pressure than when they responded at their own pace. If one assumes that the only process at work was automatic stereotyping, then such findings would lead to the strange con-

clusion that people had stronger stereotypical associations when they responded quickly. Instead, the most likely interpretation is that stereotypical representations were the same for participants responding fast and those responding slow. Other processes critical to responding accurately must have been influenced by speeded responding.

Findings like these suggest that indirect measures reflect not only automatic processes but also controlled processes. The distinction between automatic and controlled processes allows researchers to predict, for example, that when people are tired, distracted, or rushed, they are more likely to respond based on automatic impulses than when they are energetic, focused, and unhurried (Fazio & Towles-Schwen, 1999; Strack & Deutsch, 2004). The fact that many dual-process theories predict these results is both a strength and a weakness for the theories. On the one hand, predictive power is vital for a theory's value. On the other hand, dual-process models do not usually specify exactly how automatic and controlled processes relate to each other. However, in understanding conflicts between automatic and controlled influences, it is often essential to know how they are related. Formal process models have been developed to answer these more specific questions.

One such model is the process dissociation (PD) model, originally developed by Jacoby (1991) to separate implicit and explicit memory processes. Payne (2001) applied the model to separating automatic and controlled influences in a priming task, illustrating its usefulness for studying implicit social cognition. One variant of this model assumes that automatic processes drive behavior only when control over behavior fails (Jacoby, 1991). A second variant of the model assumes that automatic processes are instead dominant, and that controlled processes only drive behavior in the absence of an automatic influence (Lindsay & Jacoby, 1994). Testing how well different models describe experimental data can be used both to answer theoretical questions about how unobservable processes are related to each other and also to create quantitative estimates of those underlying processes (Payne, Jacoby, & Lambert, 2005).

A second related model is the Quad model (Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Sherman et al., 2008). Whereas the PD variants assume that either controlled or automatic processing is dominant whenever they conflict, the Quad model assumes that either kind of process can be dominant. The Quad model adds a third parameter that reflects whether automatic or controlled processes "win" when they are in

conflict and a fourth parameter for guessing biases. Bishara and Payne (2009) have recently described a unified framework for understanding PD and the Quad model, as well as other related models such as Stahl and Degner's (2007) ABC model, within a single family of models (Payne & Bishara, 2009). These models are all multinomial models, which describe a variety of unobserved processes as probabilities. Multinomial models are agnostic about temporal order, and so they cannot differentiate between processes that occur relatively early or late. However, Klauer, Voss, Schmitz, and Teige-Mocigemba (2007) proposed a diffusion model that incorporates both accuracy (as with multinomial models) and also reaction times.

Multinomial and diffusion models are exciting new directions because they offer means to test theories about the processes underlying indirect measures and simultaneously quantifying those processes. These models also offer potential for integrating insights across many different kinds of indirect measures because the same models can be applied across many different tasks to identify their common and distinct components. One of the earliest insights provided by the application of formal process models is that experimentally induced variations in indirect measurement scores can be due to different underlying mechanisms (Gawronski & Sritharan, Chapter 12, this volume). Whereas some effects have been found to reflect genuine changes in underlying associations (e.g., Sherman et al., 2008; Stewart & Payne, 2008), others stemmed from experimentally induced impairments in executive control (e.g., Govorun & Payne, 2006; Sherman et al., 2008). Given that changes in indirect measurement scores are typically interpreted as indicating changes in the underlying mental representations, formal process models provide an important means to avoid misinterpretations of empirical data and distortions in general theorizing.

WHERE IS IT GOING?

Implicit social cognition is arguably one of the liveliest and most active research areas in social psychology. The steady increase in the number of publications is just one indicator in this regard. Given the exponential growth of implicit social cognition over the past decade, we want to consider some of the emerging themes that may shape the future of implicit social cognition. As the quantum physicist Niels Bohr put it, prediction is very difficult, especially about the future. Nevertheless, in reviewing the chapters of this Handbook, we identified three themes that may guide future research and theorizing in implicit social cognition.

A first theme—one that has been recurring in different variants—is the notorious tension between person-based and situation-based approaches. This tension is reflected in early debates about whether automatically activated associations reflect mere knowledge of cultural stereotypes (Devine, 1989) or personal attitudes (Fazio et al., 1995). It is reflected in the current controversy about the role of personal versus extrapersonal associations in the IAT (Gawronski et al., 2008; Nosek & Hansen, 2008a, 2008b; Olson & Fazio, 2004; Olson et al., 2009). And it is reflected in the debate about whether responses on indirect measures reveal stable representations in memory (Fazio, 2007) or context-dependent evaluations that are constructed on the spot (Schwarz, 2007). Even though the details of these debates differ in various ways, they are all concerned with the question of whether responses on indirect measures tell us something about the individual or something about the individual's environment. Framed in this way, the conflicting positions might bring up painful memories of the person–situation debate in social and personality psychology. However, we believe that looking back at this debate can be helpful in identifying novel, integrative ways of thinking about implicit social cognition. A particularly useful direction in this regard may be social cognitive approaches to study person × situation interactions (e.g., Cervone, Shadel, & Jencius, 2001; Mischel & Shoda, 1995). One way to do this is by investigating the role of individual experiences in shaping an individual's responses in various situational contexts (e.g., Rydell & Gawronski, 2009). A complementary approach is to study the ways that dispositions lead individuals to select different situations. Even though such an approach cannot answer every question of the just-mentioned debates, it has the potential to provide a more comprehensive picture by including (1) a person's individual experiences, (2) the context in which these experiences occurred, and (3) the role of contextual cues in activating individual experiences associated with a given context. Given that research on implicit social cognition heavily relies on both individual difference (e.g., prediction of behavior; Perugini et al., Chapter 14, this volume) and experimental designs (e.g., attitude formation and change; Gawronski & Sritharan, Chapter 12, this volume), future research adopting a person × situation interaction perspective could be helpful to shed new light on some of the recurring disputes

about whether indirect measures tell us something about the person or about the person's environment.

A second trend that seems to be emerging from implicit social cognition is enhanced cross-talk and continuously shrinking boundaries between subdisciplines. The adoption of indirect measures and theoretical models from social psychology in other fields is only one example in this regard. In fact, it seems as though many social psychologists interested in basic questions have discovered applied research as an exciting context to demonstrate the power of their methods and theories. Great examples in this regard are provided by the chapters in the last section of this book, which review insights gained from implicit social cognition in the areas of health psychology (Wiers et al., Chapter 25), clinical psychology (Teachman, Cody, & Clerkin, Chapter 26), forensic psychology (Snowden & Gray, Chapter 27), consumer psychology (Perkins & Forehand, Chapter 28), and political psychology (Nosek, Graham, & Hawkins, Chapter 29). After decades of increasing compartmentalization in psychology, we see this development as an exciting trend that may ultimately lead to a more comprehensive picture of the human mind.

A final theme that we identified in several chapters in this volume is a reanalysis of the role of mental concepts (as opposed to behavioral responses) in psychological theorizing. This trend is most explicit in requests not to interpret the responses on psychological measures as direct reflections of mental concepts. After all, responses on any kind of direct or indirect measures are behavioral responses (Amodio & Mendoza, Chapter 19; De Houwer & Moors, Chapter 10; Gawronski & Sritharan, Chapter 12; Perugini et al., Chapter 14; Sherman et al., Chapter 9). To be sure, it makes sense to assume that these behavioral responses have their roots in people's minds. However, the undisputable fact that behavioral responses are mediated by mental processes does not allow researchers to equate these responses with the mental processes that presumably underlie them. A useful example to illustrate this idea is the basic notion of formal process models of indirect measures, which treat performance on indirect measures as behavioral outcomes that are jointly determined by multiple distinct processes. From the perspective of these models, the challenge for researchers is to develop theories of mental functioning that integrate the available set of behavioral observations in a parsimonious manner. Ideally, these models also imply novel predictions about behavior, so that they do not simply provide a mentalistic redescription of the available behavioral observations.

An important insight that can help to prevent circularity is the treatment of behavioral observations (e.g., performance on indirect measures) as behaviors that need to be explained rather than as mental constructs that explain behavior. From this perspective, mental constructs can still serve an important function in explaining behavior, but only to the degree that the theoretical assumptions about these constructs imply novel predictions about behavior that can be empirically scrutinized. Recent advances in implicit social cognition have already made significant steps in this direction, and the emerging concern with these issues indicates that they might become even more important in the future. Thus, even though implicit social cognition started as an inherently mentalistic approach, it clearly has the potential to make a unique contribution to the closing "decade of behavior," and we look forward to the insights it will convey in the years to come.

ACKNOWLEDGMENT

NOTE

1. For the sake of conceptual clarity, we follow De Houwer and colleagues' (2009) recommendation to use the terms *direct* and *indirect* to describe procedural characteristics of measurement procedures and the terms *implicit* and *explicit* to describe psychological features of the constructs assessed by different kinds of measurement procedures (see also De Houwer & Moors, Chapter 10, this volume).

REFERENCES

Appiah, K. A. (2008). *Experiments in ethics.* Cambridge, MA: Harvard University Press.

Arkes, H. R., & Tetlock, P. E. (2004). Attributions of implicit prejudice, or "would Jesse Jackson 'fail' the Implicit Association Test?" *Psychological Inquiry, 15,* 257–278.

Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, J. S. Nairne, I. Neath, & A. Surprenant (Eds.), *The nature of remembering: Essays in remembering Robert G. Crowder* (pp. 117–150). Washington, DC: American Psychological Association.

Banaji, M. R., & Hardin, C. D. (1996). Automatic stereotyping. *Psychological Science, 7,* 136–141.

Banaji, M. R., Nosek, B. A., & Greenwald, A. G. (2004). No place for nostalgia in science: A response to Arkes & Tetlock. *Psychological Inquiry, 15,* 279–289.

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

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

Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology, 62,* 893–912.

Bishara, A. J., & Payne, B. K. (2009). Multinomial process tree models of control and automaticity in weapon misidentification. *Journal of Experimental Social Psychology, 45,* 524–534.

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

Blanton, H., Jaccard, J., Gonzales, P. M., & Christie, C. (2006). Decoding the Implicit Association Test: Implications for criterion prediction. *Journal of Experimental Social Psychology, 42,* 192–212.

Brewer, M. B. (1988). A dual process model of impression formation. In T. K. Srull & R. S. Wyer, Jr. (Eds.), *Advances in social cognition* (Vol. 1, pp. 1–36). Hillsdale, NJ: Erlbaum.

Broadbent, D. (1971). *Decision and stress.* Oxford, UK: Academic Press.

Cervone, D., Shadel, W. D., & Jencius, S. (2001). Social-cognitive theory of personality assessment. *Personality and Social Psychology Review, 5,* 33–51.

Chaiken, S. (1980). Heuristic versus systematic processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology, 39,* 752–766.

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

Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. (2005). Separating multiple processes in implicit social cognition: The quad-model of implicit task performance. *Journal of Personality and Social Psychology, 89,* 469–487.

Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measurement: Consistency, stability, and convergent validity. *Psychological Science, 12,* 163–170.

Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals. *Journal of Personality and Social Psychology, 81,* 800–814.

De Houwer, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology, 50,* 77–85.

De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin, 135,* 347–368.

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

Doris, J. M. (2002). *Lack of character: Personality and moral behavior.* Cambridge, UK: Cambridge University Press.

Dovidio, J. F., Evans, N., & Tyler, R. B. (1986). Racial stereotypes: The contents of their cognitive representations. *Journal of Experimental Social Psychology, 22,* 22–37.

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General, 127,* 286–303.

Dunton, B. C., & Fazio, R. H. (1997). An individual difference measure of motivation to control prejudiced reactions. *Personality and Social Psychology Bulletin, 23,* 316–326.

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). San Diego, CA: Academic Press.

Fazio, R. H. (1993). Variability in the likelihood of automatic attitude activation: Data reanalysis and commentary on Bargh, Chaiken, Govender, and Pratto (1992). *Journal of Personality and Social Psychology, 64,* 753–758.

Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition, 25,* 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.

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., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology, 50,* 229–238.

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.

Fiedler, K., Messner, C., & Bluemke, M. (2006). Unresolved problems with the "I," the "A" and the "T": Logical and psychometric critique of the Implicit Association Test (IAT). *European Review of Social Psychology, 17,* 74–147.

Fiske, S. T., & Neuberg, S. L. (1990). A continuum of impression formation, from category-based to individuating processes: Influences of information and motivation on attention and interpretation. In M. P. Zanna (Ed.), *Advances in experimental social psychol-*

ceptualizing personal and extrapersonal associations. *Social and Personality Psychology Compass, 3*, 152–170.

Payne, B. K. (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, B. K. (2005). Conceptualizing control in social cognition: How executive control modulates the expression of automatic stereotyping. *Journal of Personality and Social Psychology, 89*, 488–503.

Payne, B. K., & Bishara, A. J. (2009). An integrative review of process dissociation and related models in social cognition. *European Review of Social Psychology, 20*, 272–314.

Payne, B. K., Burkley, M., & Stokes, M. B. (2008). Why do implicit and explicit attitude tests diverge?: The role of structural fit. *Journal of Personality and Social Psychology, 94*, 16–31.

Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89*, 277–293.

Payne, B. K., Jacoby, L. L., & Lambert, A. J. (2005). Attitudes as accessibility bias: Dissociating automatic and controlled components. In R. R. Hassin, J. S. Uleman, & J. A. Bargh (Eds.), *The new unconscious* (pp. 393–420). Oxford, UK: Oxford University Press.

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.

Popper, K. R. (1934). *Logic of scientific discovery*. New York: Basic Books.

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

Rothermund, K., Teige-Mocigemba, S., Gast, A., & Wentura, D. (2009). Minimizing the influence of re-coding in the IAT: The Recoding-Free Implicit Association Test (IAT-RF). *Quarterly Journal of Experimental Psychology, 62*, 84–98.

Rothermund, K., & Wentura, D. (2001). Figure-ground asymmetries in the Implicit Association Test (IAT). *Zeitschrift für Experimentelle Psychologie, 48*, 94–106.

Rydell, R. J., & Gawronski, B. (2009). I like you, I like you not: Understanding the formation of context-dependent automatic attitudes. *Cognition and Emotion, 23*, 1118–1152.

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.

Schacter, D. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 501–518.

Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition, 25*, 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.

Sherman, J. (2006). On building a better process model: It's not only how many, but which ones and by which means? *Psychological Inquiry, 17*, 173–184.

Sherman, J. W., Gawronski, B., Gonsalkorale, K., Hugenberg, K., Allen, T. J., & Groom, C. J. (2008). The self-regulation of automatic associations and behavioral impulses. *Psychological Review, 115*, 314–335.

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

Smith, E. R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4*, 108–131.

Stahl, C., & Degner, J. (2007). Assessing automatic activation of valence: A multinomial model of EAST performance. *Experimental Psychology, 54*, 99–112.

Stewart, B. D., & Payne, B. K. (2008). Bringing automatic stereotyping under control: Implementation intentions as efficient means of thought control. *Personality and Social Psychology Bulletin, 34*, 1332–1345.

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review, 8*, 220–247.

Teige-Mocigemba, S., Klauer, K. C., & Rothermund, K. (2008). Minimizing method-specific variance in the IAT: The single block IAT. *European Journal of Psychological Assessment, 24*, 237–245.

Treisman, A. (1969). Strategies and models of selective attention. *Psychological Review, 76*, 282–299.

Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review, 93*, 239–257.

Warrington, E. K., & Weiskrantz, L. (1968). A new method of testing long-term retention with special reference to amnesic patients. *Nature, 217*, 972–974.

Wigboldus, D. H. J., Holland, R. W., & van Knippenberg, A. (2004). *Single target implicit associations*. Unpublished manuscript.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review, 107*, 101–126.

Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology, 72*, 262–274.

Wittenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology, 81*, 815–827.